

DESIGN CONSIDERATIONS FOR LOW PHASE JITTER
CLOCK GENERATORS

TECHNICAL REPORT NO. SSEL-290


DISTRIBUTION STATEMENT A

Approved for Public Release
Distribution Unlimited

1998

By

Philip Sean Stetson



**SOLID-STATE
ELECTRONICS
LABORATORY**

DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCE
THE UNIVERSITY OF MICHIGAN, ANN ARBOR

19990706 092

This report has also been submitted as a dissertation in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the University of Michigan, 1998.

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 1998		3. REPORT TYPE AND DATES COVERED Technical
4. TITLE AND SUBTITLE Design Considerations for Low Phase Jitter Clock Generators			5. FUNDING NUMBERS DAAH04-94-G-0327	
6. AUTHOR(S) Philip Sean Stetson				
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(ES) University of Michigan Department of Electrical Engineering 1301 Beal Ave. Ann Arbor, MI 48109-2122			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 33790.76-EL	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This work explores the generation and propagation of phase jitter within the microprocessor clock generator. Introducing the fundamentals of phase-lock circuits, and clock generators in particular, Chapter II overviews the necessary background information required for a more in-depth analysis. Chapter III examines the concept of phase jitter, discussing its origin, its effects on a synchronous circuit, and an analytical method for calculating phase jitter. The chapter concludes by introducing a method for simulating the frequency instability of a clock generator due to phase jitter. Chapter IV is the first of three chapters discussing clock generator designs. The design described in this chapter was fabricated in Motorola's Complementary GaAs (CGaAs) process. Chapter V details the design and test of a low voltage, high frequency clock generator that exhibits low phase jitter. The advantages and disadvantages of using delay-locked loops in clock generation is explored in Chapter VI. The work concludes in Chapter VII with a series of guidelines for the design of low phase jitter clock generators for future generation microprocessors.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED		18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED
				20. LIMITATION OF ABSTRACT UL

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
TABLE OF CONTENTS	iv
LIST OF FIGURES	vi
LIST OF TABLES	xi
 INTRODUCTION	 1
 CLOCK GENERATION	 5
PLL Basics.....	5
Phase Detectors	13
EXOR	14
JKFF	15
Phase-Frequency Detector	17
Charge Pump Phase-Locked Loops	19
Loop Filters.....	21
Oscillators.....	24
 PHASE JITTER.....	 29
Phase Jitter Definition.....	29
Analytical Estimation of Phase Jitter.....	33
Phase Jitter Simulation.....	43
 CGaAs CLOCK GENERATOR.....	 48
Detailed Design.....	49
Design limitations.....	56
Charge pump saturation	56
Non-Partitioned Layout	57
Jitter measurement	58
Design summary	60
 CMOS PLL CLOCK GENERATOR	 61
Top level loop design.....	61
Loop design.....	66
Circuit design.....	69
Charge pump/Loop filter	74
Phase-frequency Detector	83

Current-Controlled Oscillator	89
V-I Converter	99
Frequency Divider	105
Bias Generator	106
Simulation and Test Results	112
CSA Phase Jitter Simulation	113
Measurement Results	116
DELAY-LOCKED LOOP CLOCK GENERATION	120
CONCLUSIONS	132
Contributions	132
CGaAs PLL Clock Generator	132
Phase Jitter	132
CMOS PLL Clock Generator	133
Delay-Locked Loop Clock Generation	133
Future Work	134
BIBLIOGRAPHY	136

LIST OF FIGURES

Figure 1.1	Basic clock generator block diagram.....	1
Figure 1.2	Microprocessor frequency versus year reported in ISSCC.....	3
Figure 2.1	Simple phase-locked loop block diagram.....	5
Figure 2.2	Example of PLL signals in the locked state.....	6
Figure 2.3	Tracking Properties of the Phase-Locked Loop.	7
Figure 2.4	PLL block diagram with annotated transfer functions.....	8
Figure 2.5	Basic RC (Passive Lag) Filter.....	8
Figure 2.6	Behavior of PLL parameters during the tracking process.....	10
Figure 2.7	PLL block diagram with supplemental frequency-detection loop.....	11
Figure 2.8	Block diagram of a frequency synthesis loop.....	12
Figure 2.9	Basic DLL topology.....	12
Figure 2.10	EXOR phase detector behavior.....	14
Figure 2.11	EXOR phase detector transfer function.....	15
Figure 2.12	Effect of asymmetric inputs on EXOR phase detector output response....	15
Figure 2.13	JKFF phase detector behavior.....	16
Figure 2.14	JKFF phase detector transfer function.....	16
Figure 2.15	Block Diagram of the phase-frequency detector.....	17
Figure 2.16	Phase-frequency detector behavior.....	18
Figure 2.17	Phase-frequency detector transfer function.....	19
Figure 2.18	Basic charge pump topology.....	20
Figure 2.19	Active-lag filter circuit schematic.....	22
Figure 2.20	Bode plots for a charge-pump PLL.....	22

Figure 2.21	Modified passive lag filter and approximate transfer function.....	23
Figure 2.22	Basic ring oscillator block diagram.....	25
Figure 2.23	Ring oscillator capacitive tuning examples.....	26
Figure 2.24	Ring oscillator resistive tuning examples.....	27
Figure 3.1	Time uncertainty represented by phase jitter.....	29
Figure 3.2	Noise transfer function of a PLL from VCO to output.....	31
Figure 3.3	Source-coupled differential pair and associated transistor noise sources.....	34
Figure 3.4	First crossing approximation.....	35
Figure 3.5	Interstage interaction.....	37
Figure 3.6	Source-coupled pair schematic.....	41
Figure 3.7	Reference voltage generator.	42
Figure 3.8	Noise simulation model.....	44
Figure 3.9	Noise spectral density for source-coupled pair delay stage.....	44
Figure 3.10	Frequency response of the source-coupled pair delay stage.....	45
Figure 3.11	Source-coupled pair RMS noise voltage versus frequency.....	46
Figure 3.12	Source-coupled pair frequency spectrum predicted by simulation.....	47
Figure 4.1	Complete block diagram of the CGaAs PLL clock generator.....	49
Figure 4.2	DCFL OR4 logic gate.....	50
Figure 4.3	Phase-frequency detector used in CGaAs PLL.....	50
Figure 4.4	Voltage source charge pump and ripple suppressing loop filter.....	51
Figure 4.5	Charge pump switch implementation.....	51
Figure 4.6	Passive lag filter with ripple suppression capacitor.....	52
Figure 4.7	VCO delay stage in CGaAs PLL.....	53

Figure 4.8	Dual ring VCO block diagram.....	53
Figure 4.9	CGaAs PLL frequency vs. control voltage.....	55
Figure 4.10	Schmoo plot of maximum frequency versus supply voltage.....	55
Figure 4.11	CGaAs PLL annotated die photo.	58
Figure 5.1	Generic charge-pump PLL block diagram.....	62
Figure 5.2	Detailed CMOS PLL clock generator block diagram.....	63
Figure 5.3	Phase margin versus divide ratio for various feed forward gain values...65	
Figure 5.4	Loop bandwidth versus divide ratio.....	66
Figure 5.5	Current steering amplifier schematic.....	72
Figure 5.6	Generic current source charge pump block diagram.....	75
Figure 5.7	Illustration of charge sharing within the charge pump.....	76
Figure 5.8	Illustration of charge injection in a charge pump.....	77
Figure 5.9	CSA charge pump schematic.....	78
Figure 5.10	Graphical illustration of CSA charge pump operation.....	78
Figure 5.11	Active loop filter implementation.....	79
Figure 5.12	Example of charge pump operation.....	81
Figure 5.13	Charge pump output current in the phase-locked state.	82
Figure 5.14	Phase-frequency detector block diagram.....	83
Figure 5.15	AOI21 CSA logic gate schematic.....	84
Figure 5.16	CSA logic gate with V_{OH} control.	84
Figure 5.17	Regulation of V_{OH} using replica feedback biasing.....	85
Figure 5.18	Example of PFD operation with V_{OH} control.	86
Figure 5.19	Voltage to current characteristic for the voltage-controlled resistor.	86

Figure 5.20	CSA logic gate sizes used in the PFD.....	87
Figure 5.21	Minimum PFD pulse width versus PFD bias current.....	88
Figure 5.22	Net pulse width versus input phase error.....	89
Figure 5.23	CSA VCO delay stage with relevant noise current sources.....	91
Figure 5.24	Interstage interaction.....	94
Figure 5.25	Ring oscillator schematic.....	97
Figure 5.26	ICO frequency versus bias current.....	97
Figure 5.27	Frequency-to-current characteristic of the ICO over V_{DD}	98
Figure 5.28	Piecewise linear ICO model.....	99
Figure 5.29	Active current mirror V-I converter.....	100
Figure 5.30	Active current mirror V-I converter DC transfer characteristic.....	101
Figure 5.31	Active current mirror V-I converter output voltage instability.....	102
Figure 5.32	Differential V-I converter schematic.....	103
Figure 5.33	Source-coupled pair V-I converter transfer characteristic.....	104
Figure 5.34	Power supply step response of the differential V-I converter.....	104
Figure 5.35	Sense-amp D-type flip-flop schematic.....	105
Figure 5.36	Frequency divider block diagram.....	106
Figure 5.37	ΔV_{BE} bias generator concept illustration.....	107
Figure 5.38	Complete kT/q bias generator schematic.....	108
Figure 5.39	Left half of the bias generator schematic.....	109
Figure 5.40	Right half of the bias generator schematic.....	110
Figure 5.41	Bias generator currents over various power supply voltages.	111
Figure 5.42	Bias generator voltages over various power supply voltages.....	111

Figure 5.43	PLL Clock generator die photo.....	112
Figure 5.44	Output noise spectral density for the CSA delay stage.....	113
Figure 5.45	RMS Noise voltage for the CSA delay stage.....	114
Figure 5.46	Simulated frequency spectrum of the CSA delay stage.....	115
Figure 5.47	Voltage-to-frequency characteristic of open loop PLL.	116
Figure 5.48	PLL output waveform.	117
Figure 6.1	Delay-locked loop clock generator block diagram.....	120
Figure 6.2	EXOR logic waveforms illustrating frequency multiplication.....	121
Figure 6.3	Logic waveforms in a DLL clock generator.....	122
Figure 6.4	CMOS Gilbert cell multiplier.....	122
Figure 6.5	Normalized phase jitter for source-coupled pair implementations.....	128
Figure 6.6	Normalized phase jitter of the DLL in comparison to the PLL.....	129
Figure 6.7	Normalized phase jitter for current-steering amplifier implementations	130
Figure 6.8	Normalized phase jitter of the DLL in comparison to the PLL.....	131
Figure 7.1	CSA dual circuit diagram	134

LIST OF TABLES

Table 3.1	Phase jitter contributions of various clock generator components [3]..	33
Table 4.1	CGaAs PLL measured results.....	54
Table 5.1	CMOS PLL Design Specifications.....	61
Table 5.2	HP-CMOS14B Level 3 HSPICE Parameters	70
Table 5.3	Measured and simulated phase jitter results.....	118

CHAPTER I

INTRODUCTION

Many applications use circuits based on phase-lock techniques. In the areas of clock recovery and frequency synthesis, circuits such as delay-locked loops and phase-locked loops predominate. The adaptability of phase-lock circuit techniques provide many benefits such as dynamic tracking, suppression of process variation, signal synchronization, and frequency multiplication.

One aspect of frequency synthesis which has gained much attention in recent years is microprocessor clock generation. During the past several years, the advancement of microprocessor frequency has far outpaced that of the system bus frequency. This trend has created the need for on-chip circuitry to multiply the system clock frequency for use by the microprocessor. Early examples of this are the Intel 486DX2 processors of the early to mid 1990's. They included "clock doublers" which provided a clock signal to the microprocessor core with a frequency twice that of the system bus. This trend has continued with the most recent Intel microprocessors. The Pentium II now offers commercially available versions running at 450 MHz, 4.5 times the 100 MHz system clock frequency.

The conventionally accepted method for performing clock multiplication for this application is to use a phase-locked loop (PLL) clock generator. The phase-locked loop accepts an input signal at a particular frequency and has the ability to produce a signal whose frequency is a multiple of this input frequency. Simple in concept, as illustrated in Figure 1.1, the phase-locked loop is a complicated and sensitive circuit. For reasons of




Figure 1.1 Basic clock generator block diagram.

cost and flexibility, it is highly desirable that the clock generator be integrated into the microprocessor. On the surface this is not an issue, as many PLL building blocks are readily realized in a commercial digital process. However, current microprocessor trends do complicate the design of the clock generator circuit.

While wholly digital PLL implementations exist, the ever increasing frequency requirements of contemporary microprocessors necessitate the use of faster, but more sensitive, analog or hybrid digital/analog designs. The increased sensitivity of these designs leads to timing instability in the clock generator output. This is particularly true in the case of high frequency microprocessors, where the digital switching creates a very noisy environment.

Another factor which complicates clock generator design is the steadily decreasing power supply voltage of contemporary microprocessors. Reduced to mitigate the power dissipation associated with high frequency operation, the low power supply voltages prevent the use of many noise tolerant circuits.

In the recent past, the primary concern in regards to the clock distribution of a microprocessor has been the clock skew across the chip. Increasing complexity and die size have compounded the layout and simulation task presented by the clock tree. The general rule of thumb is that 10% of a processor's clock period is allotted to timing issues such as clock skew. While still a significant problem, advancement of process characterization and control, parasitic extraction, and simulation techniques have enabled designers to reduce clock skew in current generation microprocessors to 60 ps - 80 ps.

While this achievement helps, rising microprocessor frequencies have revealed another design challenge. The sensitivity of the clock generation circuitry can result in timing instability of the microprocessor clock signal. This timing instability, known as phase jitter has been reported as recently as 1996 to be on the order of 150 ps [1,2,3,4].




Figure 1.2 Microprocessor frequency versus year reported in ISSCC

The timing uncertainty represented by phase jitter adds directly to the clock skew in the timing budget.

Figure 1.2 plots the frequency of microprocessors presented at the International Solid-State Circuits Conference (ISSCC) versus the year that they were presented. A best fit line through this data reflects the industry trend. A timing budget consisting of 60 ps to clock skew and 60 ps to phase jitter represents a microprocessor operating at 833 MHz. The industry trend indicates that mainstream commercial microprocessors will achieve this frequency in 2003, though high clock-rate experimental processors have already already surpassed this mark. Obviously, timing errors due to clock skew and phase jitter must be reduced to support these clock-rates.

While clock skew is a well understood problem, it remains a serious issue for next generation microprocessor designs. Research continues in areas such as parasitic extraction and efficient simulation of the clock distribution network. There is also work being done to manage the impact of clock skew by partitioning the logic between latches so that the clock skew is masked [5]. Intel designers have chosen to mitigate clock skew

in the upcoming Merced processor through the liberal distribution of clock inputs across the device [6].

Phase jitter, on the other hand, is not a well understood problem. Timing instability in the microprocessor clock signal represents a very significant portion of the timing error budget. This work explores the generation and propagation of phase jitter within the microprocessor clock generator.

Introducing the fundamentals of phase-lock circuits, and clock generators in particular, Chapter II overviews the necessary background information required for a more in-depth analysis. Chapter III examines the concept of phase jitter, discussing its origin, its effects on a synchronous circuit, and an analytical method for calculating phase jitter. The chapter concludes by introducing a method for simulating the frequency instability of a clock generator due to phase jitter. Chapter IV is the first of three chapters discussing clock generator designs. The design described in this chapter was fabricated in Motorola's Complementary GaAs (CGaAs) process. Chapter V details the design and test of a low voltage, high frequency clock generator that exhibits low phase jitter. The advantages and disadvantages of using delay-locked loops in clock generation is explored in Chapter VI. The work concludes in Chapter VII with a series of guidelines for the design of low phase jitter clock generators for future generation microprocessors.

CHAPTER II

CLOCK GENERATION

A clock generator is used to provide a multiple of the system bus frequency to the microprocessor. The clock signal provided must be stable to ensure consistent and correct operation. The phase-locked loop (PLL) is an almost ideal circuit for such an application. A PLL is capable of accepting an input signal, and producing an output signal that is matched in frequency and phase. By including a divider in the PLL design (as will be discussed later), a PLL can also provide an output signal which is in phase with the input, but at a multiple of the input signal frequency.

However, the PLL is a sensitive circuit. It often contains analog components that are very susceptible to the switching noise inherent to a high speed digital system, such as a microprocessor. Due to this, the output of a PLL clock generator actually has a time varying frequency and phase. This varying output phase uncertainty, referred to as phase jitter, is an important parameter in the design of a PLL clock generator. The following sections will discuss the design and basic operating principles of PLL's. Succeeding chapters introduce the concept of phase jitter, and detail the design of two PLL clock generators that have been designed, fabricated, and tested as part of this research.

2.1 PLL Basics

A phase-locked loop is essentially a control system utilizing a negative feedback loop to drive the input phase difference (the error signal) towards zero. Figure 2.1

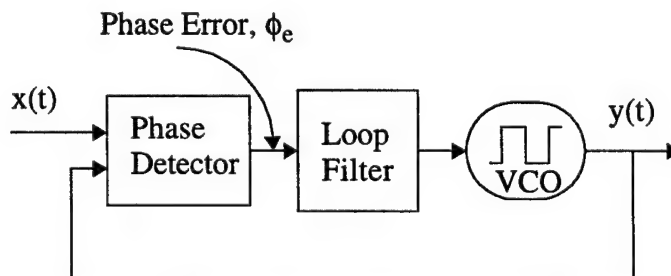


Figure 2.1 Simple phase-locked loop block diagram.

illustrated a simple PLL which consists of a phase detector (PD), loop filter (LF), and voltage-controlled oscillator (VCO).

The phase detector compares the phase difference of the two input signals, producing an output that is proportional to this difference. The phase detector acts as an error amplifier, and the negative feedback seeks to minimize this error. The loop is considered locked when the phase error is consistent, which is a result of the input and output frequencies being equal.

When the loop has reached the locked state, the PLL operates as follows. The phase detector produces a series of pulses, the width of which is proportional to the input phase difference. The loop filter smooths out the transients from this signal, producing a DC level that is proportional to the input phase difference. This DC voltage sets the frequency of the voltage-controlled oscillator. In the locked state, the VCO is biased at a frequency that is equal to the input frequency, and at some phase offset from the input signal, given by the loop dynamics. In this locked state the relevant PLL signals would look similar to those depicted in Figure 2.2.

In order to understand basic PLL operation, it is useful to examine a locked loop as it experiences a small frequency step at its input. With a slightly increased frequency, the input signal accumulates phase faster than the VCO output, which results in wider pulses at the output of the phase detector. These wider pulses produce a larger DC level at the




Figure 2.2 Example of PLL signals in the locked state.

output of the loop filter, resulting in a higher VCO operating frequency. As the frequency of the VCO increases and approaches the input frequency, the phase difference at the phase detector decreases, eventually settling at a stable value, slightly greater than before. This example, illustrated in Figure 2.3, demonstrates the tracking properties of the PLL.

Note that in the above example there were actually two processes that took place. The first was a frequency acquisition, then the loop achieved phase lock. This is an important distinction because it makes apparent two characteristics of PLL operation that must be kept in mind. First, a PLL is a system with memory. It takes a finite time for the system to react to changes at its input, and its behavior depends on its initial conditions. Second, the only means that a PLL has of correcting itself is through the VCO. Thus only changes in frequency are available. This implies that to attain phase lock, the loop may have to go out of frequency lock to accumulate phase. Eventually enough phase is acquired, attaining both frequency and phase lock. However, it is important to note that, in the locked state, the input and output frequencies are always exactly equal, though the




Figure 2.3 Tracking Properties of the Phase-Locked Loop.




Figure 2.4 PLL block diagram with annotated transfer functions.

phase error may not be, depending on the loop gain. In many applications, static phase differences are tolerable, but not even small errors in frequency can be tolerated.

To gain further insight into the operation of PLL's, one can look at the loop's transient response. However, the transient behavior of such a system is a very difficult thing to calculate, so it is more convenient to examine a linear approximation of the PLL loop dynamics. Figure 2.4 illustrates the simple PLL again with the transfer function of each block annotated within the symbol.

From this diagram one can derive the transfer function that relates the input phase to the output phase, $\frac{\phi_o}{\phi_i}(s) = H(s)$. The phase detector is represented by a subtractor with a finite gain K_{PD} . The open loop transfer function is given as $H_{ol}(s) = K_{PD}G_{LPF}\frac{K_{VCO}}{s}$. Closing the loop results in the transfer function,

$$\frac{\phi_o}{\phi_i}(s) = H(s) = \frac{K_{PD}G_{LPF}K_{VCO}}{s + K_{PD}G_{LPF}K_{VCO}}. \quad (1)$$

As the loop filter is, in its simplest form, a low pass filter, one implementation is the basic RC filter illustrated in Figure 2.5. This circuit has a transfer function given by,

$$G_{LPF}(s) = \frac{1}{1 + \frac{s}{\omega_{LPF}}}. \quad (2)$$




Figure 2.5 Basic RC (Passive Lag) Filter.

Using this relation for $G_{LPF}(s)$ results in a closed loop equation of,

$$\frac{\phi_o(s)}{\phi_i(s)} = H(s) = \frac{K_{PD}K_{VCO}}{\frac{s^2}{\omega_{LPF}^2} + s + K_{PD}K_{VCO}}. \quad (3)$$

The quantity $K_{PD}K_{VCO}$ is called the loop gain; the system is second order, gaining a single pole from each of the loop filter and VCO. By manipulating the equation somewhat, one can put it in the form of the classical second order system equation from control theory.

$$H(s) = \frac{\phi_o(s)}{\phi_i(s)} = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (4)$$

$$\omega_n = \sqrt{K_{PD}K_{VCO}\omega_{LPF}^2} \quad (5)$$

$$\zeta = \frac{1}{2} \sqrt{\frac{\omega_{LPF}}{K_{PD}K_{VCO}}} \quad (6)$$

The natural frequency, ω_n , gives an idea of the gain-bandwidth product of the loop. The damping factor, ζ , is inversely proportional to the loop gain, which presents a trade-off in PLL design.

This transfer function is essentially the response for a low pass filter, so it follows that while the loop would track and adjust accordingly for slow variation in input phase, fast variations of input phase produce only small changes at the output.

With a slowly varying input frequency ($\omega_{in} - \omega_{out} \ll \omega_{LPF}$) the loop is capable of maintaining lock as long as the parameters illustrated in Figure 2.6 vary monotonically. If the slope of and of these curves fall to zero or become negative, the loop ceases its tracking behavior. This can happen if the phase difference becomes too large and the phase detector output switches sign. In addition, the voltage controlled oscillator typically




Figure 2.6 Behavior of PLL parameters during the tracking process.

has a limited range, and its frequency to voltage gain, K_{VCO} , will fall to zero at the bounds of this range.

If the input frequency of a PLL is stepped abruptly by some amount $\Delta\omega$, the loop temporarily exits lock and the tracking behavior ensues. There is a limit on the range of $\Delta\omega$ for which the loop will regain lock. This behavior is essentially identical to the case in which a PLL is turned on and must acquire lock.

One way to look at the acquisition process is in the frequency domain. Assuming that ω_{in} is attainable by the VCO, and that the phase detector is implemented with a multiplier, the acquisition process can be explained as follows. With the VCO at ω_{fr} (the free running frequency of the oscillator) and the input frequency $\omega_{in} = \omega_{fr} + \Delta\omega$, the output of the phase detector contains a component at $\omega_{in} - \omega_{out} = \Delta\omega$. The loop filter does not completely suppress this component, so the VCO control voltage varies with frequency $\Delta\omega$. This modulates the output frequency at $\Delta\omega$ above and below ω_{fr} . When the PD multiplies the modulated component at $\omega_{fr} + \Delta\omega$, a DC component is produced at

its output. This DC component serves to drive the VCO frequency towards ω_{in} . Several cycles of such behavior may be required to drive a loop to lock. From this explanation, it should be apparent that the maximum $\Delta\omega$ depends upon how much the loop filter passes the component at $\Delta\omega$ through to the VCO. Thus, the lock range is a direct function of the loop gain at $\Delta\omega$. This suggests that the lock range of a PLL cannot be arbitrarily large because the loop gain of a PLL drops off as the difference between the input and VCO frequencies increases.

The lock range, while an important parameter, is very difficult to calculate exactly. However, since the free running frequency of a PLL can be difficult to predict, considering such factors as process variation, temperature, and other environmental factors, some assurance is needed that the PLL will operate correctly. A following section describes a design component called a phase-frequency detector, which is an alternative to the conventional multiplier. The key difference with the addition of this component, is that it allows the PLL to track in both frequency and phase, essentially extending the lock range (or, more appropriately, the pull-in range) of the loop to the limits of the VCO. An example of the pull-in range for the basic PLL discussed earlier is given in [7]. Assuming the PD is a mixer, and the LPF is the simple RC filter previously discussed, the pull-in range can be approximated as $\Delta\omega_p = \frac{\pi}{2} \sqrt{2\zeta \omega_n K_{PD} K_{VCO} - \omega_n^2}$.

Other means of extending the capture range of a PLL include the addition of a supplementary frequency detection loop as illustrated in Figure 2.7. When the frequency ω_{in} is beyond the capture range of the PLL, the frequency detection loop provides the necessary DC component at the VCO to drive the loop towards lock. As the output




Figure 2.7 PLL block diagram with supplemental frequency-detection loop.

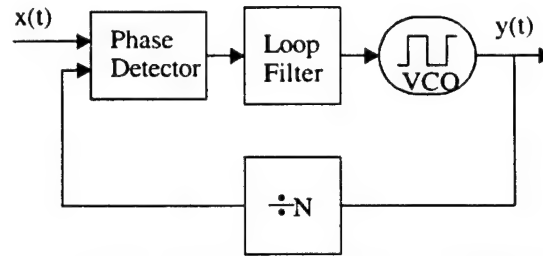


Figure 2.8 Block diagram of a frequency synthesis loop.

frequency approaches the input frequency, the contribution of the frequency loop becomes insignificant and the phase portion assumes control, driving the loop into the locked state.

The goal of clock generation is to produce a signal that is an integer multiple of the input system bus frequency. The PLL described thus far does not accomplish this goal. The addition of a divider in the feedback path, as shown in Figure 2.8, results in a PLL that operates identically to the one previously described. However, the VCO operates at a frequency which is N times that of the input frequency. Its output is divided and phase-locked to the input, ensuring that the output is a set multiple of the input frequency.

The approach illustrated by Figure 2.8 is the most straightforward for frequency synthesis. An alternative is to use a circuit called a delay-locked loop (DLL). The delay-locked loop is very similar in operation to a phase-locked loop. The DLL is most commonly used for clock recovery circuits rather than frequency synthesis, because the task of frequency division is much more readily achieved than frequency multiplication. Essentially, a DLL is a phase-lock system in which the input signal is matched to a delayed version of itself, rather than to a signal generated within the loop. The voltage-controlled delay line (VCDL) replaces the VCO and delays the input signal by a varying amount. The block diagram in Figure 2.9 represents a basic DLL. The phase detector

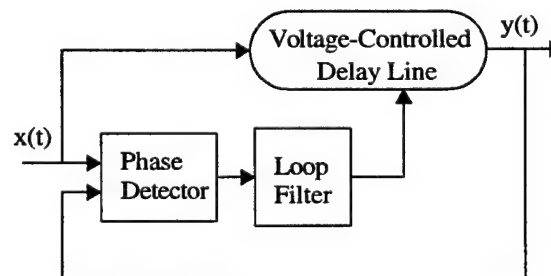


Figure 2.9 Basic DLL topology.

compares the phase of the input signal and the output of the VCDL. As in the PLL, the phase detector output is low-pass filtered to provide a near DC voltage that sets the delay of the VCDL. The negative feedback of the loop drives the phase detector towards zero phase difference. This occurs when the delay through the VCDL equals an integer multiple of the input period.

In order to use a DLL in a frequency synthesis application, it is necessary to multiply either the input signal or the output of the VCDL by the desired ratio. There are few ways to do this effectively, and all are less elegant than the PLL with frequency divider. DLLs do have two important advantages over PLL's. First, a DLL has no memory, like a PLL, and is characterized by a constant transfer function, yielding a first order open-loop response (given a first order loop filter). This characteristic gives a DLL a much more relaxed constraint between gain, bandwidth, and stability. Second, the VCDL contributes less phase jitter than a VCO. This is due to the fact that noise injected into a DLL stops at the output of the delay line, while the same noise propagates and is recycled through an oscillator in a PLL [8]. The question then, is whether the additional complexity involved in introducing frequency multiplication to a DLL overrides the gains in stability and jitter performance. Chapter VI discusses this trade-off in more detail.

The following sections overview the behavior and design of the components that make up both PLLs and DLLs. Chapter IV and Chapter V detail the design of two PLL clock generators, and expand upon the concepts introduced here.

2.2 Phase Detectors

The properties of the phase detector have a direct impact on a PLL's transient behavior, capture range, and phase-lock characteristics. The implementations vary in their phase-difference to output-voltage transfer function, their response to unequal input frequencies, and the effect that input signal amplitude and duty cycle have on their behavior. This section overviews three basic phase detector implementations.

2.2.1 EXOR

The first example is the simple exclusive-or (EXOR) logic gate. This implementation is also very similar to the linear, analog mixer and the Gilbert cell. Figure 2.10 illustrates the response of the EXOR gate to two signals that are separated by $< 90^\circ$, 90° , and $> 90^\circ$. The output of the EXOR phase detector is a signal with a frequency twice that of the inputs, and whose duty cycle is proportional to the phase difference between the inputs. When the inputs are 90° apart, as the middle plot in Figure 2.10 illustrates, the output duty cycle is 50%. This results in an average output voltage which is midway between the low and high output voltages. As the phase difference deviates from 90° , the output duty cycle ranges from 100% (when the signals are 180° out of phase), to 0% when the two signals are exactly in phase. When using an EXOR phase detector, convention denotes the 90° phase difference state as the zero, or phase-locked state because it is at this point that the average output voltage is midrange.

This yields the following response for phase deviations from -90° to $+90^\circ$, as illustrated by Figure 2.11. Assuming a rail to rail EXOR output swing, the transfer function through the EXOR gate is a constant $K_{PD} = \frac{V_{DD}}{\pi}$.

Since the EXOR gate is considered phase-locked when the input signals are 90° out of phase, the EXOR gate is often used in applications that require quadrature signals. An example of such an application is clock recovery, or data sampling, where it is desired that the sampling clock be 90° out of phase with the data. This allows the clock to sample




Figure 2.10 EXOR phase detector behavior.




Figure 2.11 EXOR phase detector transfer function.

the data when the data is most likely to be stable.

One drawback of the EXOR phase detector is that its output is dependent upon the duty cycle of its inputs. Non-symmetric input signals cause the EXOR output response to be clipped, thus reducing the overall gain through the phase detector. This problem arises from the fact that the EXOR output depends upon the *overlap* of the input signals. If the input duty cycle is less than 50%, then the reduced input pulse widths result in a similarly reduced output pulse width. This prevents the average output voltage from achieving the upper portion of its response, effectively clipping the gain at some intermediate level as depicted in Figure 2.12.

2.2.2 JKFF

The JKFF phase detector avoids the problem of input duty cycle sensitivity because it is an edge-sensitive device. This phase detector implementation operates as




Figure 2.12 Effect of asymmetric inputs on EXOR phase detector output response.




Figure 2.13 JKFF phase detector behavior.

follows. The output of the JKFF is set when the J input of the device transitions high.

Conversely, the output is cleared when the K input transitions high. Figure 2.13 illustrates this operation for the cases of $< 180^\circ$, 180° , and $> 180^\circ$ of input phase error.

Since the 180° input phase error case results in a 50% duty cycle output, it is considered the zero phase error, or phase-locked state. The input phase error ranges from 0° to 360° , the average output voltage varies from 0 to V_{DD} . This results in the transfer function as

illustrated by Figure 2.14. The transfer function through the JKFF is given by

$K_{PD} = \frac{V_{DD}}{2\pi}$. Again, this assumes that the output voltage is capable of swinging from rail to rail.




Figure 2.14 JKFF phase detector transfer function.




Figure 2.15 Block Diagram of the phase-frequency detector.

2.2.3 Phase-Frequency Detector

While the JKFF is a two-state phase detector, the phase-frequency detector (PFD) is characterized by three states. As the name suggests, this device is capable of tracking both phase and frequency. The basic logic diagram for the phase-frequency detector is shown in Figure 2.15. The PFD consists of two edge-triggered DFF's and an AND gate.

The device operates as follows. The PFD detects the edges of the input signals. Starting initially with both outputs low, a rising transition on the REF input causes the UP signal to transition high. This state indicates that the loop needs to increase its frequency in order to match the input. Likewise, when the VCO input transitions high, the DOWN output rises. This indicates that the loop needs to decrease the VCO frequency in order to match the reference input. When both outputs have switched high, the AND gate propagates a signal that resets the two flip-flops, returning the PFD to the zero state. Thus, while the circuit does reach the 1/1 state for a short period, it is suppressed by the AND gate and the PFD is essentially a three-state device. Figure 2.16 illustrates the PFD operation for three cases, $f_{\text{ref}} > f_{\text{vco}}$, $f_{\text{ref}} < f_{\text{vco}}$, and $f_{\text{ref}} = f_{\text{vco}}$ (with a finite phase error). As the logic waveforms indicate, even when the frequencies are not matched, the UP/

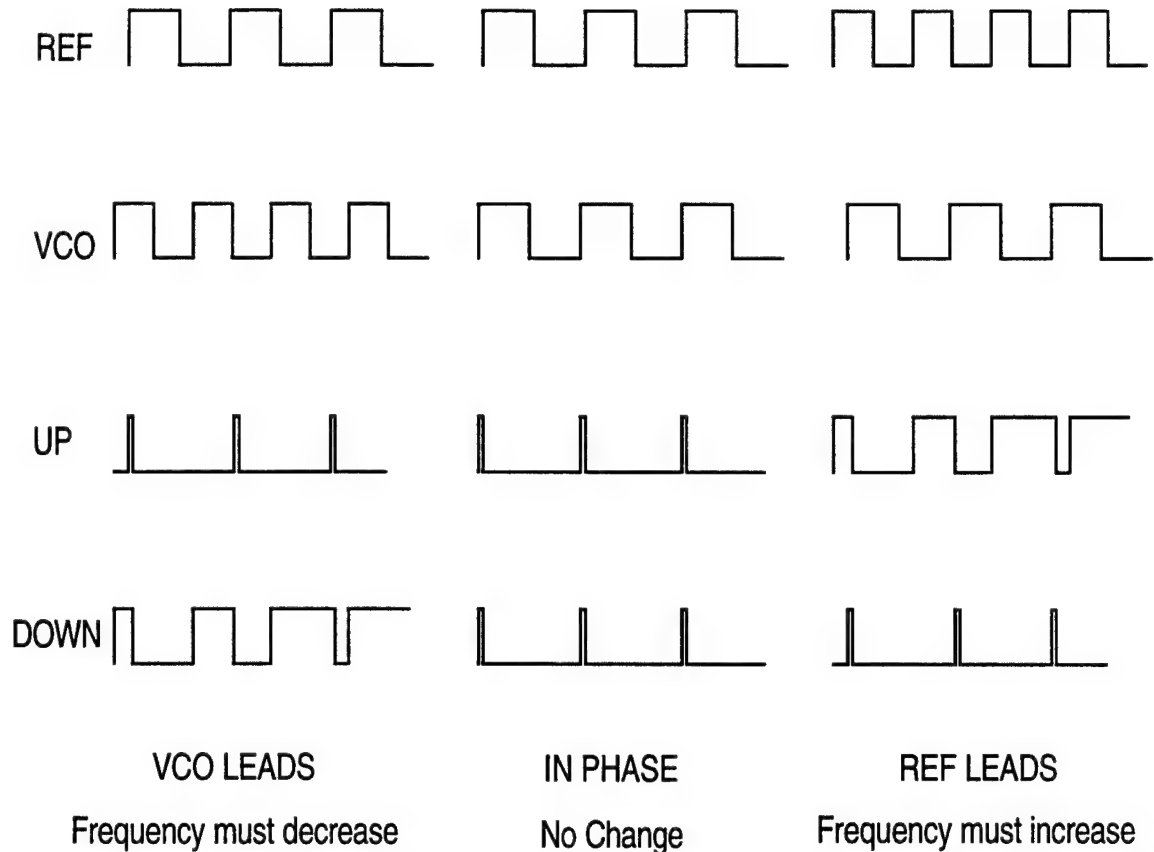


Figure 2.16 Phase-frequency detector behavior.

DOWN outputs reflect the direction in which the loop needs to be driven in order for the inputs to lock.

One difference between the PFD and the two phase detectors previously discussed, is that the PFD has two outputs, where the EXOR and JKFF phase detectors only have one. This can be handled in one of two ways. First, a differential amplifier can detect the difference between the two outputs and provide a single output which is filtered by the low-pass filter. If one looks at the average voltage between the two outputs, the response




Figure 2.17 Phase-frequency detector transfer function.

depicted in Figure 2.17 results. An important difference between this response and those of the EXOR and JKFF phase detectors is that the PFD produces an output that varies monotonically with regards to frequency error. No similar response exists for the EXOR or JKFF phase detectors. Therefore a loop using the PFD as its phase detector will lock under any condition, irrespective of the type of loop filter used. The only factor limiting the capture range of a PLL using a phase-frequency detector is the frequency range of the VCO itself.

The second method is to use a circuit known as a charge pump. The charge pump accepts both outputs from the phase-frequency detector and adds or removes charge from the loop filter in response to the pulses on UP and DOWN, respectively. The addition of the charge pump affects loop operation in several ways that will be discussed shortly.

2.3 Charge Pump Phase-Locked Loops

The use of charge pumps, and their effect on the dynamic behavior of PLLs has been well studied. The wide use of charge pumps can arguably be due to the benefits brought to PLL behavior by the phase-frequency detector. The increased tracking range




Figure 2.18 Basic charge pump topology.

and frequency-aided acquisition are two such examples, but charge pumps have their own special benefits and problems as well [9].

Figure 2.18 depicts the basic charge pump block diagram. The two switches, controlled by the UP and DOWN phase-frequency detector outputs, gate charging or discharging currents to the loop filter. Because of the phase-frequency detector operation, the two switches are only simultaneously closed for the short period of time it takes the reset path in the phase-frequency detector to propagate and clear the outputs. This results in the charge pump either sourcing current to the loop filter, sinking current from the loop filter, or presenting a high impedance node to the loop filter which prevents it from discharging on its own. If the frequency of the PLL inputs is ω_{in} rad/s, and the phase difference is denoted by θ_e , then the width of the current pulse is $t_p = \frac{|\theta_e|}{\omega_{in}}$. Given a charge pump current of I_{CP} , each corrective pulse sources, or sinks, an error current of $i_d = \frac{I_{CP}\theta_e}{2\pi}$ to, or from, the loop filter.

Given that the charge pump error current is $i_d = \frac{I_{CP}\theta_e}{2\pi}$, and denoting the loop filter transfer function as $Z_f(s)$, the VCO control voltage can be written as

$$V_c(s) = I_d(s)Z_f(s) = \frac{I_{CP}Z_f(s)\theta_e(s)}{2\pi}. \quad (7)$$

With the VCO represented by the relation K_O/s , the output phase response for a

locked loop is $\theta_o(s) = \frac{K_o V_c(s)}{s}$. Considering that $\theta_e = \theta_i - \theta_o$, the overall loop transfer functions are,

$$\frac{\theta_o(s)}{\theta_i(s)} = \frac{K_o I_{CP} Z_f(s)}{2\pi s + K_o I_{CP} Z_f(s)}, \quad (8)$$

$$\text{and } \frac{\theta_e(s)}{\theta_i(s)} = \frac{2\pi s}{2\pi s + K_o I_{CP} Z_f(s)}. \quad (9)$$

Applying the final value theorem to the phase error expression reduces it to the steady state error $\theta_{ss} = \frac{2\pi\Delta\omega}{K_o I_{CP} Z_f(0)}$ rad for a given frequency offset. Since $Z_f(0) = \infty$ for a simple passive lag filter, if this filter is implemented in a charge-pump PLL, the resulting steady-state phase error is zero. Thus, the high-impedance state of the charge pump allows a PLL using a simple passive lag filter to achieve the same results as a non-charge pump PLL using a high-gain, active filter. The charge pump permits the PLL to achieve zero static phase error without the need for DC amplification [9].

There are many different ways to implement a charge pump, and many issues that must be addressed during its design. In Section 4.2.1 and Section 5.3.1, these factors will be discussed in the context of two phase-locked loop designs.

2.4 Loop Filters

As has been shown, the choice of loop filter has a direct and significant effect on the dynamic behavior of the loop. The loop bandwidth, damping factor, and overall stability are all highly dependent upon the transfer characteristic of the loop filter, and the choice of loop filter goes hand in hand with the choice of phase detector.

If the loop does not use a charge pump, the first design decision is whether a static phase offset is acceptable. Unless a high-gain active loop filter is implemented, a non-charge pump PLL is going to lock with a finite, steady-state phase error, as demonstrated previously through use of the final value theorem. An example of a high-gain active filter




Figure 2.19 Active-lag filter circuit schematic.

is the active lag filter illustrated in Figure 2.19. For applications where a static phase offset is acceptable, however, the passive lag filter is a simple, viable choice.

For charge pump PLLs, the choice of loop filter becomes broader, but also more complex. The nearly infinite gain provided by the charge pump will drive the loop towards zero static phase error, even for loops utilizing the passive lag filter. The charge pump adds a pole at the origin to the loop transfer function. As noted above, the VCO also contributes a pole at the origin, making a charge pump PLL unstable. The bode plot in Figure 2.20 illustrates the loop's instability. To prevent this instability, one must add a zero to the transfer function to provide some phase margin. While the phase of the original design starts at -180 degrees and remains there (resulting in zero phase margin),




Figure 2.20 Bode plots for a charge-pump PLL.

the addition of the zero to the loop filter's transfer function results in a phase that rises away from the -180 degree mark. By placing the zero appropriately, one can align the resulting phase hump with the 0 dB frequency, such that acceptable phase margin exists. This results in a more stable loop topology.

Realization of a transfer function zero is accomplished using the passive lag filter by simply adding a resistor in series with the filter capacitor, as demonstrated in Figure 2.21. One problem with this implementation is that the voltage ripple at the output of the charge pump (which can be a full rail-to-rail voltage swing) gets transferred through the voltage divider of R_1 and R_2 to the VCO. This ripple causes frequency excursions that result in sideband noise about the output frequency of the loop. For this reason a capacitor is often placed in parallel with the R_2 - C_1 series combination. This second capacitor, whose pole is intentionally placed well to the right of the dominant, low frequency poles, is significantly smaller than Q . The pole placement keeps C_2 from impacting the low-frequency properties of the PLL, allowing it to behave like a second-order loop, while the filtering of C_2 significantly reduces the amount of voltage ripple that reaches the oscillator.




Figure 2.21 Modified passive lag filter and approximate transfer function.

2.5 Oscillators

Arguably the most critical component of a PLL, the oscillator is certainly the most sensitive. The oscillator is a difficult module to analyze, and simulations are often relied upon to provide insights into the various parameters of a particular oscillator design. The relevant design parameters for a voltage-controlled oscillator are the tuning range, phase jitter, supply and substrate noise rejection, and input to output characteristic linearity [8].

The tuning range of a VCO is the range of attainable frequencies for the VCO. This parameter sets the maximum and minimum output frequencies for the PLL and must be designed such that it accommodates the required input range for the PLL. The tuning range must also account for both process and temperature dependencies of the VCO.

There are two general classes of oscillators. The first is the LC-tank style oscillator. This device uses a resonant LC tank network to set the oscillator center frequency. This class of oscillator is characterized by the large, high quality inductor and capacitor structures required, and their relatively narrow tuning range. Because these structures are both costly and difficult to integrate with a digital process, and the narrow tuning range is incompatible with typical clock generator specifications, little will be said about the LC tank oscillator in this work. Perhaps in future microprocessor generations, the spectral purity offered by LC tank oscillators will outweigh their disadvantages.

The second class of oscillators includes the ring and relaxation oscillators. These are readily integrated into a digital process, and are capable of wide tuning ranges. Unfortunately, these benefits come at the cost of stability as the ring oscillator exhibits significantly less spectral purity than the LC tank oscillator. However, since the general properties of the ring oscillator match well with the requirements of a PLL microprocessor clock generator, it is worthwhile to examine their design in the context of improving their phase jitter performance.

The second and third parameters of oscillators listed above are tightly interrelated. Every oscillator is characterized by a measure of instability. This instability manifests as

phase jitter in the time domain, and phase noise in the frequency domain. The requirements of the PLL application will put an upper bound on the amount of timing inaccuracy, or spectral impurity, that is acceptable. Chapter III will provide much more detail into the origin and management of phase jitter.

The amount of phase jitter present in an oscillator is directly related to the sensitivity of the oscillator to variations in its power supply and substrate voltages. As will be shown in a following section, minimizing both the noise on the power supply voltage and the sensitivity of the VCO to this noise is absolutely critical in the design of a low phase-jitter clock generator. This task becomes much more difficult when the clock generator shares the same package and substrate with a large, high clock-rate digital microprocessor.

The frequency-to-voltage characteristic of the oscillator is important to the loop's stability. Through the tuning range of an oscillator it is desirable to have a linear response. Variation in K_O can cause nonlinearities, such as harmonic distortion, for some applications. In the context of a microprocessor clock generator, however, the VCO response can accommodate a good deal of non-linearity before stability is compromised; this nonlinearity can be as great as 30% before it becomes a problem. Though one should avoid a design in which the target frequency lies in a high-gain region, as these exhibit increased frequency instability.

The ring oscillator is implemented by chaining an odd number of inverting stages together in a ring. This topology is illustrated in Figure 2.22. It should be apparent from the illustration that the frequency of oscillation for this structure is given as $f = \frac{1}{2Nt_d}$. In this relation N represents the number of inverting stages while t_d represents the delay




Figure 2.22 Basic ring oscillator block diagram.

through a single stage with a fanout of one. As this delay is determined primarily by process device parameters, which are well characterized, the frequency of a ring oscillator is reasonably predictable.

The ring oscillator structure is readily made voltage dependent by utilizing an inverting stage, the delay of which can be varied by a voltage. Two methods are commonly employed to achieve this delay variance: capacitive tuning and resistive tuning.

Capacitive tuning uses a voltage-variable capacitor, such as a reverse-biased PN junction, or the network illustrated in Figure 2.23. This technique uses a MOS device as a means of adjusting how much of the capacitor, C , is "seen" at the output of the inverting stage. By adjusting this visible capacitive loading, the delay of the stage is varied. A drawback of capacitive tuning is that even at the minimum capacitive load, corresponding to the peak oscillator frequency, the delay stage is loaded by some additional capacitance. This additional capacitance reduces the maximum attainable frequency for a given number of delay stages. These methods are also characterized by a highly nonlinear frequency to voltage response, especially if wide tuning ranges are required.

Resistive tuning, on the other hand, provides a wide, uniform tuning range. In addition, it lends itself well to differential operation which, as will be discussed later, helps greatly to eliminate the effect of supply and substrate variation on the oscillator.




Figure 2.23 Ring oscillator capacitive tuning examples.

There are many ways to realize resistive tuning. A few of the methods are displayed in Figure 2.24.

In Figure 2.24a the load resistance of the differential stage is adjusted, varying both the time constant at the output and the small signal gain. This implementation is troubled by a few factors, the most problematic of which is that as the small signal gain decreases, the oscillation around the loop eventually dies out because the overall gain of the loop falls below unity at the frequency of operation.

The circuit of Figure 2.24b modifies the tail current of a differential stage to vary the delay. In this circuit, the small signal gain remains largely constant, but the voltage swing at the output varies. This is undesirable because when the voltage swing gets small, the oscillator is more susceptible to noise, which results in increased phase jitter.

While not a differential implementation, the delay stage shown in Figure 2.24c is commonly employed in one form or another. Known as a current starved inverter, the delay through the stage is effectively tuned through the bias on the outer transistors. The inner transistors function as a common CMOS inverter. Benefits of such an implementation are a largely constant output amplitude, the ability to operate with a lower power supply voltage, and a wide tuning range. However, the lack of a differential output,

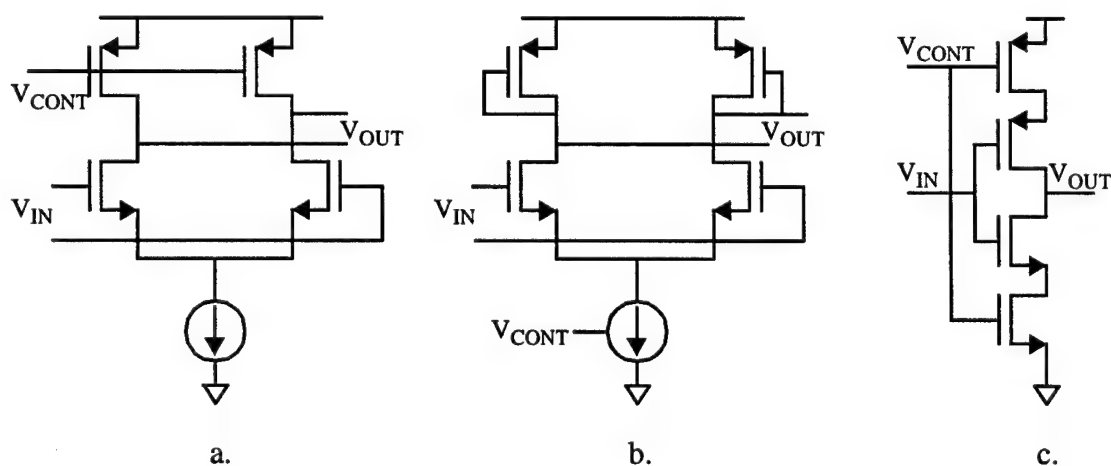


Figure 2.24 Ring oscillator resistive tuning examples.

and the rail-to-rail output swing, leaves this circuit susceptible to supply and substrate noise.

It should be noted that the number of stages in the ring is an important design decision. Fewer stages allows a ring to oscillate at a higher frequency and implies a lower power dissipation. However, the total phase shift through a ring oscillator must be 360 degrees. As the number of stages decreases, both the gain and phase shift per stage increases. Though bipolar ring oscillators have been reported that employ only two stages [10, 11], two-stage CMOS implementations typically do not operate reliably or must include additional phase shift elements that result in an oscillation frequency which is no higher than that of a three stage ring [8].

This chapter has provided the background material necessary for phase-locked loop design. Chapters IV and V will detail specific PLL clock generator designs. It is there that the trade-offs, design decisions, and analysis for the various components of a PLL clock generator are addressed.

CHAPTER III

PHASE JITTER

This chapter introduces an issue that is quickly becoming one of the forefront problems in digital system design. Increasing clock frequencies reveal phase jitter as a factor that can no longer be considered negligible. A general rule of thumb is that 10% of the clock period is allocated to clock skew [12]. When one considers that the reported clock skew figure for the DEC Alpha 21264 is 75 ps at a clock period of 1.67 ns, this doesn't appear to be a significant problem. However, as current trends indicate, microprocessor frequencies and die sizes will continue to increase. This serves to both reduce the allotted time for clock skew, and increases the difficulty of managing clock skew across the larger die. Furthermore, this timing budget has yet to consider the phase jitter of the clock distribution network, most notably within the clock generator. Since typical clock generators exhibit a peak-to-peak phase jitter on the order of 150 ps, the combination of clock skew and phase jitter becomes a very real problem as processor frequencies approach 1 GHz.

3.1 Phase Jitter Definition

Phase jitter represents the uncertainty in the sampling instant. As Figure 3.1




Figure 3.1 Time uncertainty represented by phase jitter.

shows, the actual transition time of a signal falls within some range around the nominal transition time. As a noise-related parameter, the phase jitter of a signal follows the typical statistical distribution, characterized by a mean and a variance.

The PLL clock generator is particularly susceptible because it operates on the phase of signals. Any phase jitter introduced into such a system is transferred to the output, and hence to the microprocessor or other digital system.

Phase jitter is the result of noise injected into the clock distribution network. The scope of this work deals primarily with the phase jitter of the PLL clock generator. The injected noise which results in phase jitter can come from any of these sources:

1. Noise coupled through the circuits' power supply and substrate connections.
2. Noise coupled through adjacent or intersecting traces.
3. Noise inherent to the circuits' transistors themselves.[13]

When considering phase jitter in the PLL clock generator, it is important to note that any of the loop components can contribute phase jitter to the PLL [14]. A necessary step is to understand how the phase jitter (phase noise when considered in the frequency domain) is propagated to the output.

To evaluate the contribution of noise at the input of a PLL, one considers the classic second order response. If the input is a pure sinusoid with an excess input phase

$$\frac{\phi_{OUT}}{\phi_{IN}} = H(s) = \frac{\omega_n^2}{s^2 + 2\zeta \omega_n s + \omega_n^2}$$

that does not vary with time, then $s = 0$ and the output $H(s) = 1$. Similarly, with a very slowly varying input phase, the transfer function remains very close to unity and the tracking properties of the PLL function as expected. However, as the input phase variation increases in frequency, the low pass filter properties of the PLL become apparent. The excess output phase will drop and eventually approach zero. The input phase noise transfer function is shaped by the low pass filter characteristics of the PLL, as

represented by the bandwidth of the loop. It is this property of PLLs that inspired their predominant use in the communications industry. Thus, to minimize the output phase jitter, in response to phase variation at the input, the loop bandwidth should be minimized. This creates a trade-off, however, as decreasing the bandwidth decreases stability and increases the lock time and capture range of the PLL.

Since, for monolithic implementations, the VCO is the primary source of phase jitter when compared to the other loop components, it is very useful to examine the transfer of VCO phase jitter to the output. If the VCO phase jitter is modeled as an additive component as illustrated in Figure 3.2, and a strictly periodic signal is applied to the input, the transfer of VCO phase to the output phase is given by the following relation.

$$\frac{\phi_{OUT}(s)}{\phi_{VCO}(s)} = \frac{s(s + \omega_{LPF})}{s^2 + 2\zeta \omega_n s + \omega_n^2}$$

This relation was the same characteristic equation as given above, but the transfer of VCO phase jitter to the output also contains two zeros. These zeros at $s = 0$ and $s = -\omega_{LPF}$ mean that the characteristic of this function is that of a high pass filter. This result makes sense because the zero at $s = 0$ implies that slowly varying phase jitter at ϕ_{VCO} results in a very small output phase jitter. This is expected because a slowly varying ϕ_{VCO} gives the loop time to propagate this phase difference through the phase detector, which will produce a DC output into the VCO that opposes the phase difference caused by ϕ_{VCO} . However, as the frequency of ϕ_{VCO} increases, the loop eventually becomes unable to




Figure 3.2 Noise transfer function of a PLL from VCO to output.

correct for it, and ϕ_{VCO} is passed directly to the output phase, ϕ_{OUT} .

This effect is demonstrated well by a common case when testing the noise immunity of PLLs. The test involves applying a small step to the PLL power supply voltage and observing the time required for the PLL to settle out the resultant input/output phase difference [15]. As will be discussed later, a step on the power supply voltage predominantly affects the VCO. Thus, the transfer function from ϕ_{VCO} to ϕ_{OUT} can be employed to get a first order approximation of the PLL response to a step on the power supply voltage. Presuming that the voltage step produces a phase step of ϕ_1 , the following relation represents the output response.

$$\phi_{OUT}(t) = \phi_1 e^{-\zeta \omega_n t} \left[\cos \sqrt{1 - \zeta^2} \omega_n t + \frac{\zeta}{\sqrt{1 - \zeta^2}} \sin \sqrt{1 - \zeta^2} \omega_n t \right] \quad [8]$$

It is apparent from this relation that the output phase jumps to ϕ_1 in response to the input step, and then oscillates, decaying towards zero with a time constant $(\zeta \omega_n)^{-1}$. This implies that a designer should maximize the quantity $(\zeta \omega_n)$ for fast recovery of the PLL.

These two cases present conflicting design needs. In order to minimize the transfer of phase jitter from the input to the output, it is desirable to have a small loop bandwidth. However, to allow the PLL to recover quickly from phase step disturbances in the VCO, it is desirable to maximize the loop bandwidth. Usually, this design trade-off can be largely resolved by considering the target application of the loop. In the case of a PLL clock generator, the input clock is typically a stable signal from a crystal oscillator which has little phase jitter, while the PLL is integrated into a noisy digital environment where it is subjected to power supply steps and other noise. In this application, it is obvious that a PLL clock generator should seek to maximize the loop bandwidth. Of course, as will be discussed later, it is also desirable to minimize the PLL's sensitivity to such noise, thus reducing the initial phase error represented by ϕ_1 in the previous example.

3.2 Analytical Estimation of Phase Jitter

Of the three sources of noise in a digital system, power supply and substrate, interconnect-coupled, and intrinsic device noise, this work will primarily address the effects of the first and third. While a constantly present issue, the coupling of signals to the clock line is typically addressed through the use of guard traces to shield the clock line. Further examination of the issues associated with cross talk, as CMOS process metallization pitches grow more and more dense, is beyond the scope of this work.

In 1996, VonKaenel et al. reported a PLL design for which they showed both measured and simulated data for the phase jitter contribution of the various PLL components. These contributions were evaluated both in a clean environment and in the presence of power supply and substrate voltage noise. Table 3.1 repeats these findings.

Jitter contributor without supply noise	P-P Phase Jitter (ps)
White Noise in VCO	30
Dead zone of PFD	<10
Leakage on LF and Charge injection	15
Total Jitter without supply noise	55
Jitter due to a 0.2 V supply jump in 30 ps	
VCO induced jitter	80
Jitter induced by the change of the LF voltage	10
Total Jitter due to a 0.2 V supply jump	90
Jitter due to a 10 mV substrate jump in 30 ps	
VCO induced jitter	<5
Total Jitter due to a 10 mV substrate jump	5
Total Jitter (sum of the above contributors)	150

Table 3.1 Phase jitter contributions of various clock generator components [3]

The white noise of the VCO refers to the noise generated by the transistors that compose the VCO. This inherent transistor noise, composed of thermal, shot, and flicker noise, is unavoidable. Though it has previously been insignificant in the context of a microprocessor clock generator, rising clock-rates continue to make it a more immediate issue. In fact, if the magnitude for this inherent noise given in Table 3.1 is noted, it is




Figure 3.3 Source-coupled differential pair and associated transistor noise sources.

apparent that even this quantity will become a significant contributor to the overall phase jitter. For this reason, a method is needed to estimate the jitter due to the inherent transistor noise in the VCO.

Such an analytical method, based on a VCO implemented as a ring oscillator with source-coupled-pair differential inverting stages, is presented in [16]. The schematic in Figure 3.3 shows the source-coupled pair, complete with its inherent transistor noise sources. This method is applicable to other circuits too, as will be demonstrated in Section 5.3.3. Since a ring oscillator using this delay stage operates at a frequency determined by the number of stages and the delay of each stage, the effect of these inherent noise sources on the stage delay must be evaluated.

To begin the analysis, it is assumed that each VCO stage contributes a delay denoted by t_d . The delay has a timing error due to the transistor noise, denoted by Δt . This error is defined to have a mean value of zero and a variance given by $\overline{\Delta t^2}$.

Considering that each stage drives some load capacitance C_L , with a total differential output swing given by $2V_{PP}$, the delay through a single stage is approximated by $t_d = V_{PP} \frac{C_L}{I_{SS}}$. This assumes that the next stage begins switching when the differential

output of the previous stage reaches zero. The quantity I_{SS} represents the tail bias current of the source-coupled differential pair, and the factor $(\frac{C_L}{I_{SS}})$ represents the slew-rate.

The first crossing approximation estimates the timing error as illustrated in Figure 3.4. The first crossing approximation makes the simplifying assumption that the next stage will begin switching when its input crosses some nominal threshold [17]. In this case, that threshold is the point where the differential inputs cross zero. As this figure depicts, an error voltage at the nominal time of crossing will result in a timing error. This timing error, whose magnitude is proportional to the magnitude of the error voltage and the signal slew rate, delays the actual time of crossing. Thus, the timing error variance is given by equation (8).

$$\overline{\Delta t^2} \equiv \overline{\Delta v_n^2} \left(\frac{C_L}{I_{SS}} \right)^2 \quad (10)$$

The noise voltage, Δv_n , is the sum of the contributions from the noise sources depicted in Figure 3.3. At this point, the analysis can take two different paths. It is simpler to assume that the noise voltage is equivalent to that when the circuit is in equilibrium. However, it is more accurate to consider the time varying behavior of the




Figure 3.4 First crossing approximation.

noise sources. As in [16], both methods will be shown for completeness. This provides a means of assessing the quality of the equilibrium approximation.

If the noise sources are taken to be equivalent to the equilibrium values, traditional noise analysis techniques apply, as demonstrated in [17]. The output noise voltage can then be determined by integrating the noise spectral density over the bandwidth of the low pass filter formed by the output load resistor and the input capacitance of the next stage. Applying this result to the equations for t_d and $\overline{\Delta\tau^2}$ yields the following expression for the single-stage RMS timing jitter normalized to the delay through the stage.

$$\frac{\Delta\tau_{rms}}{t_d} \cong \frac{\Delta v_{rms}}{V_{PP}} = \sqrt{\frac{2kT}{C_L}} \left(\sqrt{1 + \frac{2}{3}a_v} \right) \frac{1}{V_{PP}}$$

The term $\sqrt{1 + \frac{2}{3}a_v}$ is referred to as the noise contribution factor and is denoted by ξ . The first component represents the PMOS load device, and the second represents the NMOS driver device. The second component depends upon the voltage gain since, for a fixed output bandwidth, a higher gain implies a higher transconductance and thus a higher noise contribution.

To consider the time varying behavior of the noise sources, one must determine when each of the sources contributes significantly. This is done in a qualitative manner by considering the regions of operation for a source-coupled differential pair: balanced and unbalanced. The tail current noise generator is assumed zero in the balanced mode and at full contribution for the unbalanced region. Conversely, the NMOS driver noise generator is considered full and constant for the balanced region, and zero for the unbalanced. The PMOS load device noise generator constantly contributes to the output noise through both

regions. Combining these contributions, and using autocorrelation and convolution, results in the following relation for the noise contribution factor as given in [16].

$$\xi = \sqrt{1 + \frac{2}{3}a_v \left(1 - e^{-\frac{t}{\tau}}\right) + \frac{2\sqrt{2}}{3}a_v e^{-\frac{t}{\tau}}}$$

The second term represents the source-coupled differential pair contribution which rises steadily from time $t = 0$ to the value given in the time-invariant analysis. The third term shows the decay of the tail current source contribution from its peak value in the unbalanced state. The time constant is approximately the delay through the gate, t_d , so the exponentials essentially reduce to constants at $t = t_d$, and the noise contribution factor depends primarily upon the gain a_v .

The time varying behavior of the noise sources is not the only second order effect that should be considered. It has been assumed thus far that the noise in one stage affects only the switching behavior of the following stage. In fact, a chain of inverters often exhibits overlapping transitions, where more than one successive inverter is in the active amplifying region at any particular time. Therefore, a more accurate model would consider the noise contributions from the previous two stages.

This interstage interaction is calculated by considering two successive stages, and determining the voltage noise at the output of the second stage which originates from the thermal noise sources in the first stage. Figure 3.5 illustrates this concept, and expands the




Figure 3.5 Interstage interaction.

idea to show the circuit model for the interaction. Solving for v_{n2} yields a slightly different noise contribution factor ξ and an increase in the voltage noise variance by a factor of $\frac{1}{2}a_v^2$. Therefore, the new normalized timing jitter is given by,

$$\frac{\Delta\tau_{1rms}}{t_d} = \sqrt{\frac{kT}{C_L}} \frac{1}{V_{GS} - V_T} \xi.$$

The analysis thus far has determined the timing jitter for a single stage; we want to derive the cycle-to-cycle phase jitter for a ring oscillator. For an N-stage ring oscillator, the oscillator period is given by $2N \times t_d$, and the total phase jitter variance for one cycle of operation, will be $2N \times \overline{\Delta\tau_1^2}$. Using the normalized jitter from the previous analysis, the total jitter variance is given by the following expression.

$$\overline{\Delta\tau_N^2} = \overline{\Delta\tau_1^2} \times \frac{T_o}{t_d} = \frac{kT}{I_{SS}} \frac{a_v \xi^2}{V_{GS} - V_T} \times T_o \quad (11)$$

From this relation, it is apparent that with everything else fixed, the ring oscillator jitter improves linearly with power supply current. This establishes a power dissipation-phase jitter performance trade-off. However, the result only necessarily holds for the class of circuits considered in this analysis. A useful step is to write the total jitter, normalized to the oscillation period, T_o . When this is done, the phase jitter variance is shown to have a $\frac{1}{\sqrt{T_o}}$ dependence. Thus, oscillators with higher frequency have worse jitter performance.

This result applies only to a stand alone voltage-controlled oscillator. When such an oscillator is incorporated into a phase-locked loop, an important effect occurs. The ring oscillator timing jitter was derived for a single cycle of operation. In actual operation, timing errors experienced in one cycle will propagate through and be compounded by timing errors in the next cycle. In a PLL implementation, the action of the PLL will

attempt to correct this trend, but its effectiveness is dependent upon the bandwidth of the PLL.

The work reported in [18] shows that the accumulated phase jitter will include the timing error of all those cycles for which the PLL has yet to take corrective action. The PLL transfer function for a typical second order charge pump PLL is examined for a series of phase steps with random magnitude. By summing the PLL responses to these phase steps it is shown that the jitter variance of a PLL is given by $\Delta\tau_{PLL} \approx \alpha \frac{2\pi\Delta\tau_{rms}}{T_o}$, where $\Delta\tau_{rms}$ is the $\overline{\Delta\tau_N^2}$ for a ring oscillator and $\alpha = \sqrt{\frac{1}{2K_dK_waT_o}}$. This α is defined as the accumulation factor, and it is inversely proportional to the square root of the quantity $K = K_dK_wa$, which is the bandwidth of the PLL. The accumulation factor is typically in the range of 10-100 [16]. It is interesting to note that for clock recovery applications, where the input source is characterized by significant jitter, a small bandwidth PLL is desired to suppress this input jitter. However, in frequency synthesis applications, the input source is typically very clean, and the PLL bandwidth should be made as large as possible to suppress the phase jitter produced by the VCO.

A similar analysis done for a delay-locked loop has shown that its accumulation factor is equal to one [18]. This conclusion is obvious when the action of a delay line is considered. In a voltage-controlled delay line, the individual stages contribute the same amount of jitter as an identical stage in a ring oscillator, but the propagation of the jitter ends at the end of the delay line, rather than being compounded cycle after cycle, as it is in a ring oscillator. This difference means that the jitter of a DLL is significantly lower than a similarly designed PLL.

Returning to Table 3.1 from [3], the remaining supply-independent jitter contributions (the dead zone of the phase-frequency detector, loop filter leakage and charge injection) are very implementation specific. The phase-frequency detector dead zone is defined as the minimum input phase difference that results in an output response. Any dead zone in an implementation obviously translates directly into phase jitter, as the

loop will allow itself to slip that far out of lock without taking corrective action. Dead zone was traditionally a problem because the charge pump was located off-chip, forcing the phase-detector UP/DOWN outputs to drive large capacitive package and PCB loads. For small input phase differences, the reset path through the PFD would clear the UP/DOWN outputs before they were able to propagate to valid levels, resulting in a dead zone on the order of ns of magnitude. The standard solution is to introduce delay in the reset path to allow the outputs to reach valid voltage levels. In contemporary designs, however, the charge pump is on-chip so the PFD dead zone is rarely a problem.

Nonlinearities in the loop filter and charge pump, such as leakage currents, charge injection, and charge sharing can contribute significantly to both a PLL's jitter performance and overall dynamic behavior. These issues will be discussed in Section 5.3.

The rest of the phase jitter components listed in Table 3.1 are the result of noise, either on the power supply or substrate. These contributions comprise the majority of the PLL phase jitter. Thus, power supply independence of the PLL, and particularly of the VCO, is critical with regard to minimizing the total PLL phase jitter.

A widely accepted method for minimizing the effect of supply and substrate noise is to implement both the control voltage and oscillator paths through a PLL with completely differential circuits. The common mode rejection of such circuits is well understood, and this property results in attenuation of any power supply noise by a factor known as the power supply rejection ratio (PSRR), which is typically on the order of 70 - 90 dB.

Implementing both the control voltage path and oscillator in differential logic applies this rejection property to the most sensitive portions of a PLL. It must be noted, however, that implementations with complementary signals which swing rail-to-rail cannot be considered truly differential because the rail-to-rail signals exhibit poor power supply rejection [8]. Differential implementations are characterized by two main problems. First, differential circuits dissipate quiescent current. A three-stage ring




Figure 3.6 Source-coupled pair schematic.

oscillator, having a load capacitance of 100fF on each stage, will need to sink nearly 1mA of current per stage to oscillate at 1GHz. The expression $I = C \frac{\Delta V}{\Delta t}$ provides a first order means of estimating the time required to drive the output capacitance. While many microprocessors operate at power levels high enough to render the power dissipation of the clock generator insignificant, a growing number of applications require low power processors.

Perhaps the more important issue concerning the use of differential circuits is the voltage headroom required for their proper operation. The simple source-coupled pair, illustrated in Figure 3.6, requires adequate biasing for three transistors stacked between the power supply and ground rails. These transistors must be kept in the saturation region to retain the circuit's rejection properties. As power supply voltages are continuously lowered to reduce power dissipation, it becomes increasingly difficult to design a differential ring oscillator circuit. Limited output swing (which is more susceptible to phase jitter), and relatively high current levels for fast operation, also make implementing differential circuits infeasible for contemporary digital processes.

Another method for suppressing the phase jitter due to power supply noise is to prevent the power supply step from being seen by the sensitive circuitry in the first place.




Figure 3.7 Reference voltage generator.

This technique was demonstrated in [19]. The circuit diagram in Figure 3.7 depicts a feedback network which establishes a floating voltage which is locked to some reference voltage (0.5 V in this case). The sensitive circuitry is placed between V_{DD} and this floating voltage which is used to sink current like the ground node. The benefit of such a scheme appears when one considers how this 0.5 V reference voltage is generated. A bandgap reference generates the constant current I_0 . Since a bandgap reference is largely power supply independent, the voltage dropped across the resistor R will remain constant. Thus, the voltage $V_{ref} = V_{DD} - I_0 R$. The variance of V_{ref} with respect to V_{DD} is given by

$$\frac{\delta V_{ref}}{\delta V_{DD}} = 1 - R \frac{\delta I_0}{\delta V_{DD}} = 1 .$$

Thus, the reference voltage, V_{ref} , varies identically with V_{DD} . The voltage across the resistor is the difference between V_{DD} and V_{ref} , which, from the above expression, does not vary with power supply voltage. This maintains a constant voltage across the sensitive circuitry, isolating it from power supply variation.

Since the jitter induced by a step on the supply voltage is proportional to the magnitude of the voltage step [3], this circuit reduces jitter by a factor equal to its ability to track supply variation. This ability depends primarily on the degree to which the current

I_0 remains constant with power supply voltage. The performance of this circuit is very implementation specific, but [19] reports a worst case power supply sensitivity of 4.7 %/V, a four-fold improvement over the sensitivity of the oscillator without reference-voltage circuitry. A side benefit of this method is that temperature independence can be achieved if the temperature coefficient of the voltage reference cancels out the temperature dependence of the oscillator.

An obvious drawback of this technique is that it also requires differential circuitry. Both the generation of the power supply independent current, and the floating reference voltage require voltage headroom. This is not as significant a problem as it was in the case of the oscillator, however, because the required current levels and dynamic range are significantly lower. A distinct benefit of such a technique is that it allows the designer to use non-differential techniques for the sensitive circuits, passing the noise rejection requirements on to the reference generator.

3.3 Phase Jitter Simulation

Incorporation of a phase jitter model in the simulation of phase-locked systems is essential. While the analytical method described previously allows one to estimate the phase jitter in an oscillator delay stage, it assumes a linear scaling of single-stage phase jitter to oscillator phase jitter. Simulation of the oscillator jitter based on that of a single stage would be a more accurate approach. Additionally, such a method provides the simulation framework for including the oscillator phase jitter in system-level simulations.

The method proposed in [16] derived the intrinsic transistor noise, and inferred from that an estimate of timing jitter through the first crossing approximation. It follows that simulating the phase jitter in a ring oscillator should also start with a circuit noise analysis.

The model in Figure 3.8 represents the circuit on which the noise analysis is performed. Inverting stages #1 and #2 are DC biased such that both are simultaneously




Figure 3.8 Noise simulation model

transparent, as they would be during switching. This insures that any interstage interaction is included within the simulation model. A source-coupled pair delay stage is used to illustrate the proposed jitter simulation methodology. This makes it possible to compare the results with those predicted by the analytical model given by equation (11).

The noise analysis results in the noise spectral density given in Figure 3.9. It represents the frequency composition of the noise at the differential output of the second stage. Given in units of $\frac{V}{\sqrt{Hz}}$, it must be integrated over the circuit bandwidth to derive the RMS noise voltage. Performing a frequency response simulation on the source-coupled pair circuit provides the gain versus frequency plot illustrated by Figure 3.10. The bandwidth is chosen to be the frequency at which the gain is 3dB below the DC level. Squaring the noise spectral density, integrating, and taking the square root of the result




Figure 3.9 Noise spectral density for source-coupled pair delay stage




Figure 3.10 Frequency response of the source-coupled pair delay stage

provides the RMS noise voltage. The RMS noise voltage, as a function of frequency, is shown in Figure 3.11. At the 3dB bandwidth, the RMS noise voltage evaluates to 135.4 μV . In the analytical method, the first crossing approximation calculates a timing jitter from this voltage based on the output slew rate. This introduces inaccuracy because the first crossing approximation makes assumptions about the circuit's slew rate and switching behavior [17,18,].

While the noise distribution is not solely given by white noise, at high frequencies the shot noise and thermal noise effects dominate and can be accurately modeled with a white noise generator. To include such a white noise generator into a SPICE-like simulation environment (Analogy's Saber in this case), a pseudo-white noise generator must be used, since true white noise models are not available.

White noise is defined by a constant power spectral density over frequency. Matlab includes a good band-limited white noise generator primitive. Generating a noise vector of suitable length, and applying each element of this vector to the timesteps of a

piecewise-linear voltage source, effectively forms a white noise generator. This provides a source of white noise appropriate for use in transient circuit simulations.

The piecewise linear voltage source is readily incorporated into the oscillator circuit as a source in series between the output and input of each successive ring oscillator stage. The sources, while having identical RMS voltages, are differentiated by using different random seeds in Matlab to generate the white noise vector. Simulating the oscillator over many periods and measuring the frequency distribution results in the frequency spectrum displayed in Figure 3.12.

The target source-coupled pair ring oscillator, biased at a tail current of 500 mA, oscillated at a nominal frequency of 627.4 MHz. The frequency spectrum, resulting from the noise simulation, exhibited 0.5 MHz of spread. This represents a peak cycle-to-cycle phase jitter of 0.508 ps. For the same source-coupled pair oscillator, equation (11) predicts a peak cycle-to-cycle phase jitter of 0.182 ps. Thus, the simulation method predicts the same order of magnitude as the analytical method. Furthermore, the




Figure 3.11 Source-coupled pair RMS noise voltage versus frequency

simulation method made no assumptions about the delay stage, or the effects of noise upon it. An added benefit to this method is that it can be used in system-level PLL simulations to predict the effects of inherent oscillator noise on the tracking characteristics of a PLL design.

This chapter has established the need for stable clock generator circuits and has demonstrated ways in which noise in the system, and within the generator itself, causes the timing instability known as phase jitter. Design trade-offs and techniques for minimizing phase jitter were introduced. The chapter also described both analytical and simulation methods for predicting inherent oscillator phase jitter. The simulation methodology provides a more accurate prediction of the phase jitter due to the inherent transistor noise by eliminating the assumptions made in the analytical derivation.




Figure 3.12 Source-coupled pair frequency spectrum predicted by simulation

CHAPTER IV

CGaAs CLOCK GENERATOR

One of the primary concerns in high frequency microprocessor design is power dissipation. Since dynamic power dissipation follows the fundamental relation $P = fCV^2$, it is apparent that power dissipation is proportional to frequency. It follows that increases in microprocessor frequency cause increases in power dissipation. However, the more significant term in the relation is the quadratic dependence upon power supply voltage. This implies that the increase in power, due to increasing frequency, can be offset by a decrease in power supply voltage. This is, in fact, the primary strategy adopted by the computing industry to manage the power dissipation in microprocessor designs. The clock generator discussed here is designed to operate on a supply voltage of 0.9 to 1.5 V. This is comparable to what is predicted for mainstream CMOS microprocessors in upcoming generations.

The intended system is a multi-chip, PowerPC-based microprocessor with a target clock rate of 1 GHz and a system clock rate of 100 MHz. The following PLL design is intended for application within the MCM as a global clock generator, providing a high-frequency processor clock in-phase with the system clock.

The targeted technology for the clock generator is Motorola's Complementary Gallium-Arsenide (CGaAs) process. This technology is a three layer metal process that implements both n-type and p-type heterostructure devices and lends itself well to a variety of logic styles. The gate metal is also usable as a local interconnect layer and can be patterned to widths of 0.5 μ m.

4.1 Detailed Design

The PLL operates in a typical fashion. The components, as illustrated in the block diagram of Figure 4.1, are implemented with digital circuits (except for the analog low pass filter). Operation can be summarized as follows.

The phase detector compares an external clock signal input with the output of the divide-by-N counter. The phase detector then produces a series of pulses which represent the phase difference between the signals at its inputs. These pulses drive a charge pump that slowly injects or removes charge from the low pass filter's capacitor. The voltage across the capacitor is essentially the control voltage for the voltage-controlled oscillator. The charge pump then decreases or increases this voltage which varies the VCO frequency accordingly. Clock drivers buffer the VCO's output to produce the processor clock. The processor clock is also fed into the divide-by-N counter where it is divided in frequency, by some integer N, and fed back into the phase detector. This system forms a negative feedback loop that, when designed correctly, reaches a steady state condition of nearly zero phase error. In this "phase-locked" state, the rising edge of the processor clock is synchronized with the rising edge of the system clock.

The logic style used to implement the digital blocks of the PLL is a Direct-Coupled FET Logic (DCFL). This DCFL style realizes the logic function with n-type devices and employs a single p-type pull-up device with its gate tied to ground. The

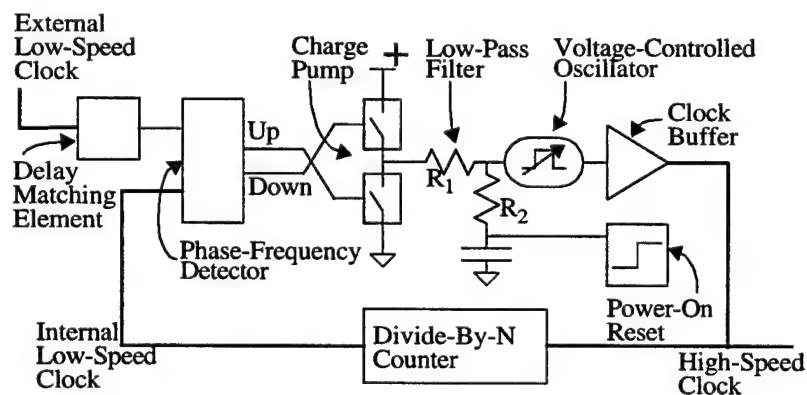


Figure 4.1 Complete block diagram of the CGaAs PLL clock generator.




Figure 4.2 DCFL OR4 logic gate.

schematic in Figure 4.2 demonstrates the DCFL style. Fast gate speed, transistor efficient implementations, and a static power dissipation characterize the DCFL logic style.

The three-state phase-frequency detector was chosen for the many properties discussed earlier. The implementation was adapted from one used in [15]. As Figure 4.3 demonstrates, the NAND gates of the original implementation were replaced by NOR gates because NOR gates are faster in the DCFL logic style used. The delay through the PFD is important, as it determines the minimum width of the output pulses when there is zero input phase difference. This width should be small, but nonzero, as these simultaneous pulses ensure that there is no dead zone in the PFD implementation.

The choice of low pass filter design affects the dynamic performance of the PLL. Since the PFD, when used in combination with a charge pump, enables the PLL to achieve lock at any input frequency attainable by the VCO, the simple passive lag filter can be




Figure 4.3 Phase-frequency detector used in CGaAs PLL.




Figure 4.4 Voltage source charge pump and ripple suppressing loop filter.

used and will achieve similar PLL performance to a more complex filter implementation.

Figure 4.4 illustrates the schematic for the charge-pump and low pass filter block.

The switch is implemented with the circuit depicted in Figure 4.5. Two complementary pass gates control the application of V_{DD} or ground to the charge pump output. Inverters and simple delay gates ensure that both sides of the complementary pass gates open simultaneously.

The use of a voltage charge pump such as this one results in a rail-to-rail voltage swing at the output of the charge pump. If allowed to feed directly into the control input of the VCO, the frequency excursions of the VCO would be $\Delta\omega = K_{VCO}(V_{DD} - V_C)$ or $\Delta\omega = K_{VCO}V_C$, depending upon the state of the charge pump (where K_{VCO} represents the frequency-to-voltage gain of the oscillator, and V_C is the control voltage). For many




Figure 4.5 Charge pump switch implementation.




Figure 4.6 Passive lag filter with ripple suppression capacitor.

applications, including frequency synthesis, these excursions result in an unacceptable level of spectral impurity. For this reason, the resistor R_2 is added to the simple passive lag filter. The sizes of R_1 and R_2 are chosen such that the majority of the voltage excursion falls across R_1 , thus reducing the magnitude of the frequency jumps. An alternative that was not implemented is to use a two pole filter as illustrated in Figure 4.6. The addition of the parallel capacitor introduces filtering which helps mitigate the voltage ripple. In this design, the value of R_1 was chosen to be ten times that of R_2 . Since the loop should theoretically lock to $\theta_e = 0$, the frequency excursions should be eliminated in the locked state.

The low pass filter is completely integrated. It uses n-type diffusion to form the resistors $R_1 = 47.2 \text{ k}\Omega$ and $R_2 = 4.72 \text{ k}\Omega$. A 100pF capacitor is implemented as a four-layer stack. This stack, which measures roughly $800\mu\text{m} \times 800\mu\text{m}$, dominates the die area.

The voltage-controlled oscillator was implemented using an adaptation of the configuration reported in [1]. As illustrated in Figure 4.7 the frequency variation is achieved through the use of both current-starved and variable-capacitance tuning methods. The VCO ring oscillator inverting stage is essentially a DCFL inverter whose pull-up p-type device is split. One part has its gate tied to ground, like the DCFL gate in Figure 4.2, while the second part's gate is tied to the control voltage, providing a variable pull-up strength.

This configuration produces a pull-up slew rate that is voltage dependent. A drawback to this method is that the duty cycle of the oscillator waveform is also made




Figure 4.7 VCO delay stage in CGaAs PLL.

voltage dependent. This was mitigated, as depicted in Figure 4.8, by the use of two ring oscillators with cross-coupled inverters between the two rings. The combined rings structure produces a complementary clock signal that is very close to a 50% duty cycle. One should also note the p-type devices connected between the two ring oscillators across the stages. These devices ensure that the two ring oscillators power up 180 degrees out of phase [1].

The ring oscillators were designed with three stages. The voltage variable RC delay network at the output of each inverting stage is a voltage controlled resistor and 4-




Figure 4.8 Dual ring VCO block diagram.

layer metal stack capacitor. The voltage-controlled resistor provides a more linear capacitive load-to-voltage response than does a single transistor [15].

The combination of current-starved inverter stages and output capacitive tuning utilizes the entire range of control voltage. The VCO frequency to voltage relationship has no regions for which the VCO has zero gain, or is inoperable. This results in increased stability and dynamic range.

The divide-by-N counter consists of a 4-bit down counter that toggles the output of a fifth flip-flop every N cycles of the VCO clock. Four control bits set the divide ratio from 2 to 16. This output is the low-frequency clock that completes the negative feedback loop at the phase-frequency detector.

The design of the down counter is basic and not covered here in detail. It should be noted, however, that the output flip-flop (that toggles every N clock cycles) is synchronized with the VCO output clock. This reduces the static phase error between the input system clock and the output processor clock to a single clock-to-Q delay of a D-type flip-flop. To further reduce this static error, the input system clock is used to toggle an identical D-type flip-flop. The output of this flip-flop is then passed on to the PFD, thus reducing the steady state error to transistor variation and interconnect delay.

The design was fabricated and tested. The following table and figures illustrate the measured PLL system performance. Table 4.1 summarizes the operating characteristics measured at a power supply voltage of 1.5 V and an input clock frequency of 100 MHz.

Table 4.1 CGaAs PLL measured results

Maximum VCO Frequency	775 MHz
Minimum VCO Frequency	137 MHz
Peak-to-Peak Jitter	< 120 ps
Lock-in Time	6.7 μ s
Power Dissipation	300 mW @ 1.5 V
PLL Die Area	1.4 mm ²

The graph in Figure 4.9 displays the PLL frequency-to-voltage relationship. The data was measured by sweeping the VCO control voltage from an external test point and




Figure 4.9 CGaAs PLL frequency vs. control voltage.

observing the resultant open-loop frequency. As expected, the combination of current-starved inverters and capacitive tuning provided a gradual frequency response over the entire voltage range. Furthermore, the cross-coupled inverters nearly eliminated the duty-cycle voltage dependence. A worst-case (control voltage approaching V_{DD}) duty-cycle of 50.2 % was observed.

The schmo diagram of maximum VCO frequency versus power supply voltage is shown in Figure 4.10. This diagram illustrates the wide functional voltage range, reaching as low as 0.8 V.

The lock-in time parameter was measured using a 100 MHz frequency step between various operating ranges. Within the nominal operating range of the PLL, the worst case lock-in time was observed to be 6.7 μ s.




Figure 4.10 Schmo plot of maximum frequency versus supply voltage.

4.2 Design limitations

The testing of this design revealed several design issues that limited both the PLL's functionality and testability. This section will detail these findings as a means of setting the stage for the next PLL clock generator design.

4.2.1 Charge pump saturation

The most significant problem occurs in the charge pump. While the VCO in this design is capable of an output frequency of nearly 800 MHz, the closed loop system fails to lock at output frequencies above 550 MHz. There are two likely explanations for this discrepancy.

The first involves the charge pump implementation. As illustrated in Figure 4.4, the charge pump is implemented as a three-state switch which applies either V_D , ground, or a high impedance to its output node. The VCO is designed such that lower control voltages produce higher frequencies. Therefore, charge pump operation must be examined at low control voltages.

As the loop drives the charge pump output voltage toward ground, the voltage across the lower switch (S1) decreases accordingly. This reduced drain-source voltage results in lower discharge current drive. Since the PLL is trying to drive its output frequency higher, the UP output from the PFD will be active a higher percentage of the time than the DOWN output. While this may be the case, the voltage across S1 will eventually fall below the saturation voltage of the switch transistors. Once this occurs, the switches begin operating in the linear region and their current drive decreases with their drain-to-source voltage. At some point the amount of charge removed from the loop filter with each DOWN pulse will become comparable to the charge added with the short UP pulse which occurs each cycle. This results in a lowering of the loop gain, and the action

of the negative feedback ceases. Thus, the loop fails to achieve a locked condition at the high end of attainable VCO frequencies.

The second possibility involves the divide-by-N counter. Since the VCO is often designed to operate at the upper frequency limits for a given technology, the divide-by-N counter is commonly strained to meet that frequency. If the VCO is capable of producing frequencies that are beyond those at which the divide-by-N can operate, it is possible that the closed loop will be broken, as no divider output will reach the phase-frequency detector. The lack of a VCO-derived input at the PFD is interpreted as the case where the VCO must increase in frequency to match the reference input. This results in a runaway loop that is driven towards the upper end of the VCO frequency range. The scenario is unlikely in this particular case, however, as the observation of a saturated output frequency, below the VCO maximum, suggests the former explanation.

4.2.2 Non-Partitioned Layout

Another design issue, which became apparent during testing, concerns the PLL layout. The PLL consists of several blocks. The die photo in Figure 4.11 illustrates the modular layout of these blocks. However, the photo also shows the large buffers used to distribute the high-frequency output clock signal to the output drivers. One of these buffers is even used to drive the signal off of the chip. Unfortunately, the layout was implemented in such a way that these buffers use the same power distribution traces as the PLL core blocks. This is a problem for a couple of reasons.

First, the buffers draw a lot of current and certainly contribute a significant amount of switching noise to the power rails. As has been discussed, power supply noise is a major contributor to overall phase jitter. For this reason, sensitive components, notably the VCO, should be isolated from such noise. Furthermore, these sensitive components




Figure 4.11 CGaAs PLL annotated die photo.

should have decoupling capacitance local to the block. These oversights undoubtedly resulted in increased phase jitter.

The second problem associated with the layout is the fact that an accurate measurement of the PLL power dissipation is unattainable. The high-power buffers sourced from the PLL core power rails skew any measurements. Since these PLL's are often used in systems where power dissipation is critical, such as portable computers, accurate monitoring of a PLL's power dissipation is necessary to evaluate the quality of a particular implementation.

4.2.3 Jitter measurement

Phase jitter is a difficult quantity to accurately measure in the laboratory. This should be apparent simply from the fact that the parameter of interest must be measured to picosecond accuracies. Furthermore, the metric that is most meaningful is cycle-to-cycle

jitter. This requires that the test environment have a stable, clean signal at the frequency of interest to serve as a temporal reference. The parameter measured in this design was absolute phase jitter, measured with a high frequency digital oscilloscope with its persistence set to infinity. The resultant spread of the VCO output clock edges reveals the stability of the clock signal in an absolute sense. Unfortunately, the absolute phase jitter is not as meaningful a quantity as the cycle-to-cycle phase jitter.

To measure the cycle-to-cycle phase jitter, it is necessary to compare the period of one cycle to the previous cycle. The difference between the two is the cycle-to-cycle phase jitter. By measuring this quantity over a large number of cycles, during which noise is injected into the system, the peak cycle-to-cycle phase jitter can be obtained.

One technique for measuring cycle-to-cycle phase jitter involves a delay coil. A delay coil is simply a long length of wire compactly coiled, which is designed to have a particular delay from input to output. By feeding the PLL output signal into the coil, and tuning the PLL so that the nominal output period matches the coil delay, the input and output of the coil can be simultaneously observed on an oscilloscope. By using the coil output as a trigger signal for its input, the previous clock cycle becomes the reference for the succeeding clock cycle. Observing the resulting spread with this setup reveals the cycle-to-cycle phase jitter. The drawback to this method is that the delay coil is only useful at a single frequency.

Another technique involves using an external signal as the measurement reference. By setting this external reference signal to the nominal output frequency of the PLL (for a particular measurement), the variation of the PLL output period can be observed through the same configuration as the delay coil method. The period variation represents the cycle-to-cycle phase jitter. This method is limited by the availability of a precise, high-frequency signal generator.

A more flexible technique involves post-processing of the measurement data. If a high sampling rate, digital oscilloscope is used to observe the PLL output signal, the

sampled waveforms can be saved to disk. While this method requires the availability of an oscilloscope with the capability of saving data via disk, or a test interface such as HPIB, the post processing step can readily produce a wide range of measurement results, including peak cycle-to-cycle phase jitter. As with the previous two methods, this one is limited by the sampling rate of the oscilloscope. Additionally, the technique is constrained by the storage capacity of the media used to record the measurement data, as the postprocessing step requires a large number of cycles.

4.3 Design summary

The CGaAs PLL discussed in this section, while relatively simple in design, demonstrated, for the first time, the feasibility of such a device in this technology. Though layout decisions limited its performance, the observed performance surpassed that of contemporary CMOS designs. The speed and low power supply voltage of CGaAs makes it a very attractive technology for such circuits, many of which are used in low power, portable applications. The primary drawback of using CGaAs in frequency synthesis designs is the lack of accurate modeling available for the process. HSPICE models adequately predict DC behavior, as the transistor active I-V characteristics are well modeled. However, the leakage currents and node capacitances of the CGaAs transistors are not predicted accurately by current HSPICE models. This poses a significant problem for detailed analog circuit design in general, and especially for frequency synthesis circuits. The problem is exacerbated by the relative immaturity and process instability of CGaAs. Since the process is still under development, lot-to-lot parameter variation is often significant. This variation can easily lead to reduced performance or inoperability in analog circuits. However, it should be noted that none of these issues are fundamental ones. They are all related to the immaturity of CGaAs as a design process. Fundamentally, the process is both sound and attractive for such circuits. This is even more true when one considers the exceptional radiation hardness exhibited by CGaAs.

CHAPTER V

CMOS PLL CLOCK GENERATOR

The testing of the CGaAs PLL revealed several limitations of the design. These limitations, coupled with the goal of exploring new techniques to reduce the PLL phase jitter and the need for a more stable process, prompted the design of a second phase-locked loop clock generator. This second design was implemented in Hewlett-Packard's 0.5 μ m CMOS process. This process was, at the time, the best available through the MOSIS fabrication service.

The goal of the CMOS PLL design was to achieve next generation output frequency, jitter performance, and power supply voltage, using current generation technology. Table 5.1 summarizes the desired design specifications. The following sections will detail the design, layout, and test of the CMOS PLL clock generator.

Table 5.1 CMOS PLL Design Specifications

Specification	Design Goals
Power Dissipation	< 15 mW
Frequency Range	400 - 1000 MHz
Divide Ratio	2 - 32
Peak Cycle-to-Cycle Jitter	< 50 ps

5.1 Top level loop design

The general topology of a phase-locked loop clock generator is fairly consistent across designs. Particular designs are differentiated more by the implementation of the individual blocks, than by differences in topology. The goal of the top level loop design is to ensure stability of the loop's dynamic behavior across a range of operating conditions.

At this point, the specifics of the circuits are not known, but the results of this analysis provide the guidelines to which they will be designed.

Stability requires that the loop have sufficient gain and phase margin across the range of possible loop parameters. The general block diagram for a charge pump PLL is shown in Figure 5.1. The two poles at the origin, produced by the charge pump and the VCO, require the addition of a zero in the loop filter for stability. This zero is most readily realized in the loop filter transfer function by putting a resistor in series with the filter capacitor in a passive lag implementation. With such an implementation, the zero location can be set appropriately through the choice of resistor and capacitor values. The method suffers from the problems inherent in using resistors in a standard silicon process.

First, the process variation in both diffusion and polysilicon resistors is significant. While ratios can be accurately predicted, absolute values can vary by up to 20%. This adds another dimension of variation that affects the PLL's stability. Second, if relatively accurate metal resistors are used, the implementation becomes prohibitively large (die area relates directly to the cost of the device). While laser trimming can be used to adjust the zero location, this is an expensive process that is generally avoided. Making all these issues more problematic is the fact that such an implementation provides no tunability to compensate for either process variation or loop parameter variation, such as the possible divide ratio settings. The divide ratio directly impacts the loop dynamic behavior by

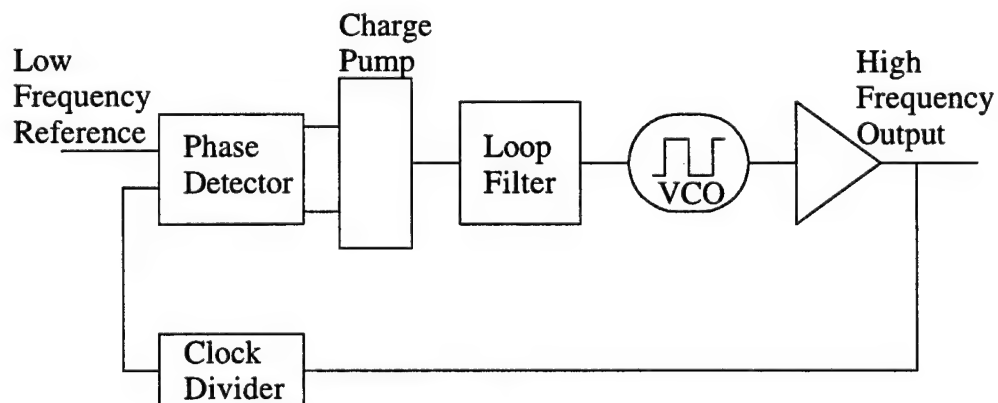


Figure 5.1 Generic charge-pump PLL block diagram.

dividing the VCO gain, K_{VCO} . If a wide range of divide ratios is desired, the loop must be proven stable for all potential values. This can lead to design decisions that may compromise the performance of some configurations, in order to retain stability over the whole range.

One way to avoid this problem is to examine what the control voltage in such a system looks like. The voltage across the filter $V_C(s) = i(s) \times R + \frac{i(s)}{sC}$, where $i(s)$ is the portion of the charge pump current that flows through the series R-C combination. If this voltage is then used to generate a bias current for a resistively tuned voltage-controlled oscillator, the resultant bias current, $i_C(s) = i(s) \times RA_g + \frac{i(s)}{sC} A_g$. The factor A_g represents the transconductance of the voltage-to-current converter. It is apparent that the total control signal is a combination of the charge pump current multiplied by some scaling factor, and the integral of the injected charge from the charge pump.

Figure 5.2 illustrates a variation of a loop topology reported in [20], in which a feed forward path is added to represent the zero in the overall loop transfer function. This topology utilizes a current-controlled oscillator and a voltage-to-current converter, to translate the output of the loop filter into an oscillator bias current. The loop also uses a second charge pump, called an auxiliary charge pump, whose output adds directly to the oscillator bias current. Thus, the oscillator bias current $i_C(s) = i(s) \times A_f + \frac{i(s)}{sC} A_g$. This is the same form as that using a passive lag filter with the zero implemented by a resistor.

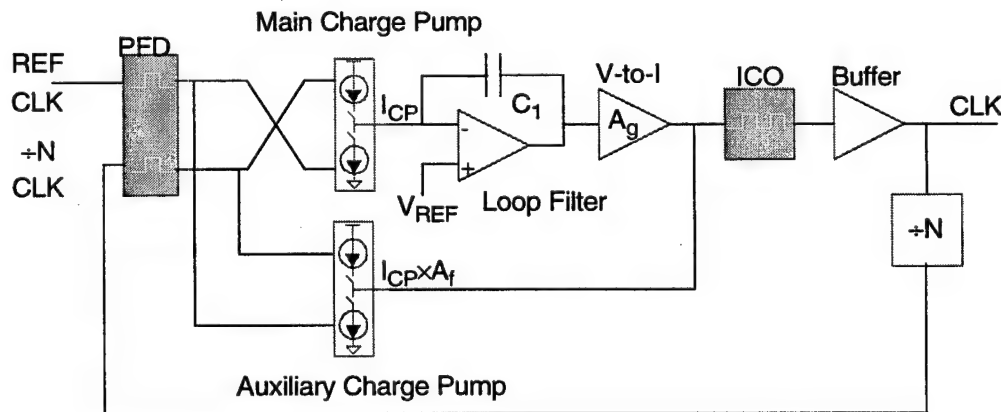


Figure 5.2 Detailed CMOS PLL clock generator block diagram.

The feed forward path adds a zero to the transfer function without the need of resistors, thus eliminating the aforementioned problems associated with realizing resistors in a digital CMOS process.

The open loop transfer function equation for this topology is,

$$G(s)H(s) = \frac{I_{CP}K_{ICO}}{N} A_g \frac{1 + \frac{A_f}{A_g} s C_1}{s^2 C_1 (1 + T_i s)} . \quad (12)$$

In this loop design, the filter is implemented with an active integrator, giving the loop filter transfer function of $\frac{1}{s C_1}$. The pole represented by the $(1 + T_i s)$ term is produced by the intrinsic capacitance of the oscillator bias input. The pole can be tuned by adding capacitance to this input. As the open loop transfer relation shows, the loop contains a zero represented by the numerator term $1 + \frac{A_f}{A_g} s C_1$. Moreover, the location of this zero can be changed by varying the ratio between the feed forward current gain, A_f , and the voltage to current transconductance, A_g . This ratio is readily tunable, providing flexibility in the loop's dynamic behavior.

The open loop transfer function can be evaluated to see the effect that various loop parameters have on the overall loop stability. These analyses also demonstrate the control over stability that is gained through the feed forward implementation. For example,




Figure 5.3 Phase margin versus divide ratio for various feed forward gain values.

Figure 5.3 shows the phase margin for various values of the feed forward gain, versus the divide ratio, N . For stability, a phase margin of 40-50 degrees is desired.

While there are values for A_f that attain this phase margin across the target range of N , the bandwidth for these configurations is shown by Figure 5.4 to be too small for robust loop dynamic behavior. The discussion in Section 3.2 established that one factor in the VCO phase jitter in a PLL is inversely proportional to the square root of the PLL bandwidth. For this reason, it is desirable to have the flexibility to tune the PLL bandwidth for the particular operating region, rather than designing a loop that is singularly stable over all operating conditions. The latter case simply puts unnecessary constraints upon the design. The feed forward implementation provides the means for the PLL to achieve better performance over all modes of operation.




Figure 5.4 Loop bandwidth versus divide ratio.

5.2 Loop design

The PLL clock generator was simulated at the block level using the behavioral simulation capabilities of Analog's Saber simulation tool[22]. Its mixed-mode simulation environment provides the ability to simulate at the transfer function level, and systematically work towards a full circuit-level simulation as the design progresses.

First order design of the system starts by assuming that the loop behavior is dominated by the poles at the origin, while the pole due to the intrinsic oscillator input capacitance has a less significant impact. This assumption allows the system to be discussed as a classic second order system, with the bandwidth and damping factor given

by equations (13) and (14), respectively. The assumption is valid as long as the PLL bandwidth is closer to the location of the zero than the third order pole [23].

$$\omega_n = \sqrt{\frac{K_{ICO} I_{CP} A_g}{2\pi C_1}} \quad (13)$$

$$\xi = \frac{\frac{A_g}{A_f} C_1}{2} \sqrt{\frac{K_{ICO} I_{CP} A_g}{2\pi C_1}} \quad (14)$$

The design specifications shown in Table 5.1 allow the calculation of preliminary values for the loop parameters. To a first order, the bias current range required to achieve the 1200 MHz to 400 MHz frequency range is given by $I = C \frac{\Delta V}{\Delta t}$.

$$1.0 \text{ GHz} \rightarrow \Delta t \approx 200 \text{ ps} \quad I = (200 \text{ fF}) \frac{1 \text{ V}}{200 \text{ ps}} = 1 \text{ mA} \quad (15)$$

$$400 \text{ MHz} \rightarrow \Delta t \approx 600 \text{ ps} \quad I = (200 \text{ fF}) \frac{1 \text{ V}}{600 \text{ ps}} = 0.33 \text{ mA} \quad (16)$$

$$\therefore K_{ICO} = \frac{1000 \text{ MHz} - 400 \text{ MHz}}{1 \text{ mA} - 0.33 \text{ mA}} = (8.96 \times 10^{11}) \times 2\pi \quad \frac{\text{rad}}{\text{s} \cdot \text{A}} \quad (17)$$

Assuming that the active integrator and voltage-to-current converter are capable of operating across 90% of the available voltage range, the transconductance of the voltage-to-current converter is approximately,

$$A_g = \frac{1 \text{ mA} - 0.33 \text{ mA}}{1.8 \text{ V} - 0 \text{ V}} = 0.37 \quad \frac{\text{mA}}{\text{V}}$$

The charge pump current, I_{CP} , is a parameter that has some flexibility in its choice. The value must be large enough to be well above the noise inherent to such circuits, but must be small enough that each corrective pulse only affects the VCO control voltage by an incremental amount. Too large a value causes voltage deviations at the VCO large enough to overdrive the VCO and cause erratic loop behavior. Typical values for I_{CP} are in the 10's of μA . A preliminary value of 15mA is chosen for this design.

The filter capacitor is also a parameter that gives the designer a measure of freedom. The size of the filter capacitor directly affects the stability of the loop. This is

apparent looking at the simple relation $\Delta V = \frac{\Delta Q}{C}$. The charge pump will inject a $\Delta Q = I_{CP}t_d$, where t_d represents the pulse width from the phase-frequency detector (which is proportional to the input phase error). This ΔQ will affect the voltage by an amount inversely proportional to the filter capacitance. If the capacitance value chosen is small, each ΔQ will have a large effect on the control voltage. This is not desirable for loop stability, as it tends to produce large excursions of the VCO frequency, magnifying any mismatches or non-idealities in the control path design. Thus, loop filter capacitors are often made large, within the bounds allowed by the particular application. Since die area translates directly to cost, it is necessary to evaluate this trade-off carefully. Besides cost, another effect of filter capacitor size is the lock time of the loop. The dynamic behavior of a loop with a large filter capacitor is characterized by a slow response time. This is reflected in the fact that the loop bandwidth is inversely proportional to the square root of the loop filter capacitance. In clock generator applications, the lock time is often an unimportant factor (there are some exceptions), but the slow response time and low bandwidth have another consequence.

Remembering the analysis of Section 3.2, the phase jitter generated in the VCO of a PLL is actually amplified by a factor α , which is inversely proportional to the loop bandwidth. Thus, the jitter performance improves with decreasing loop filter capacitance, but can also worsen due to mismatches and non-idealities in the charge pump. This arises because like deviations in current cause larger voltage deviations across smaller capacitors. The optimal point in this trade-off is unclear, but appears to be very design specific. It weighs the phase jitter inherent in the VCO against the phase jitter produced in response to non-idealities in the charge pump and control path implementation. The analysis and simulation in Section 3.2 and Section 3.3 predicted inherent phase jitter numbers on the order of 5 ps, given a PLL with a 1 MHz bandwidth. In this design, a relatively large loop filter capacitance of $C_f = 400\text{pF}$ was chosen to mitigate the effect of

charge pump non-linearities, with the intent of minimizing phase jitter through careful design of the oscillator and control path.

With preliminary values established for all of the loop parameters, it is possible to analyze the stability of the system for various values of A_f . The curves in Figure 5.3 illustrate the phase margin for various loop configurations. Both N and A_f affect the PLL dynamic behavior. The curves also show that values of A_f from 2 to 8 provide the flexibility required to produce a well-conditioned loop at any of the desired settings for N . Furthermore, this flexibility provides the means of compensating for parameter changes due to process variation.

The realization of the variable A_f is a circuit implementation issue which will be addressed in a following section. Armed with these preliminary PLL loop parameters, circuit implementation for the various blocks can commence. The next section deals with general circuit issues. The subsequent sections detail the design of each individual block. With this background, the whole circuit and its layout will be discussed. The chapter concludes with a discussion of the measured PLL characteristics.

5.3 Circuit design

While each component of the phase-locked loop must be designed separately, the goal for each remains the same, to minimize the phase jitter. To a large extent, this means minimizing the components' sensitivity to power supply noise. In some cases, such as the oscillator, there are other considerations as well.

A generally accepted method for attaining insensitivity to power supply variations is the use of differential logic. The source-coupled differential pair illustrated in Figure 3.6 is commonly used as a delay stage in PLL and DLL designs [27]. Reasons for this

include the fully differential signalling, high frequency operation, and good power supply noise rejection.

In this design, with power supply, frequency, and jitter specifications that are demanding in a 0.5 μm process, differential logic cannot be used, as will be shown in the following example. Table 5.2 lists the minimum, average, and maximum threshold voltage and square-law current gain for the n-type and p-type transistors in the HP 0.5 μm process. These numbers were compiled from parametric data sheets on the MOSIS web site. Using these values, the DC biasing for a differential stage can be calculated.

Table 5.2 HP-CMOS14B Level 3 HSPICE Parameters

		Max	Min	Mean	Std Dev
NMOS	VT0	0.7118	0.6566	0.6722	0.016
	KP ($\times 10^{-4}$)	1.9647	1.6885	1.7878	0.0696
PMOS	VT0	-0.95	-0.8887	-0.9275	0.0178
	KP ($\times 10^{-5}$)	4.874	3.8312	4.3096	0.2933

Using the voltage-controlled oscillator as an example, the biasing requirements must first be determined. The oscillator is intended to operate at frequencies as high as 1.0GHz. As calculated previously, this implies an approximate bias current of 1mA. Since phase jitter in a source-coupled differential delay stage is inversely proportional to the voltage swing, the voltage swing should be maximized [16]. This value was chosen to be 50% of the available voltage swing, or 1.0 V. The source-coupled pair (SCP) of Figure 3.6 (on p. 41) requires biasing such that the devices remain in saturation to ensure proper operation.

$$\text{Assume } V_{DSAT0} = 0.3V \rightarrow I_0 = \frac{1}{2}K_{pmin}\left(\frac{W}{L}\right)_0 (V_{GS0} - V_{T0max})^2 \Rightarrow \left(\frac{W}{L}\right)_0 = 132$$

To provide voltage margin so that the device is biased solidly in saturation, the drain-source voltage of the current source transistor, M_0 , should remain greater than 0.5V.

The load transistors, M_3 and M_4 of Figure 3.6, require a similar calculation to determine the transistor sizes.

$$\text{Assume } V_{DSAT3} = 0.3V \rightarrow I_3 = \frac{1}{2}K_{pmin}\left(\frac{W}{L}\right)_3 (V_{GS3} - V_{T3max})^2 \Rightarrow \left(\frac{W}{L}\right)_3 = 580$$

To maintain a V_{DSAT} of 0.3V across the load transistors when all of the tail current is flowing through one side of the differential pair, the output voltage cannot rise above $2.0V - 0.95V - 0.3V = 0.75V$. This implies that only 0.25V remains as V_{DS} for the input transistor. To keep the input device saturated in this case requires $\frac{W}{L} \geq 525$. This input transistor size produces a load capacitance that is far in excess of the 200fF which was assumed, and this number does not include the drain capacitance of the output devices in the previous stage. Thus, in order to maintain the desired Δt , either the tail current must increase or the voltage swing must decrease. The tail current cannot increase, for the previous analysis has just shown that the current requirements were initially too large. Since lower voltage swings are more susceptible to phase jitter [8], reducing the voltage swing is not an attractive option.

This example reveals that at the tail current levels required to achieve oscillation frequencies above 1GHz, the target technology does not support SCP differential implementations. There is simply too much voltage headroom required to properly bias three devices stacked between the power and ground rails.

To a certain extent, this situation was precipitated by the use of a 0.5 μ m process (with threshold voltages intended for 3.3V operation) for a 2.0V application. However, to maintain acceptable noise margins, there is a limit to how low the threshold voltages can be made. Thus, future generations of CMOS could see such a situation arise. In fact, Dennis Buss, a Texas Instruments fellow and vice-president in charge of analog mixed-signal development, was recently quoted in EE times, stating that “the headroom’ for traditional analog circuits, like amplifiers, is lost. There is little margin between the




Figure 5.5 Current steering amplifier schematic.

amplitude peaks of drivers and the current noise of a typical CMOS circuit.” [28]. A solution to this problem would prove quite useful.

One possible approach to the solution is to reduce the number of transistors between the power and ground rails. The circuit illustrated in Figure 5.5 is called a current-steering amplifier [29]. The circuit is biased with a current through the current source transistor, M_0 . Depending upon the input voltage, this current is steered between one of the two legs, producing a low output when the input is high, and a high output when the input is low.

When the input voltage is low, the drive transistor M_1 is in the cutoff region and all the current, I_0 , flows through M_2 . The size of M_2 sets the output high voltage. Since the load transistor is diode connected, V_{OH} is essentially the gate voltage required to sink the current I_0 , as given by equation (18),

$$V_{OH} = V_T + \sqrt{\frac{2I_0}{K_P \left(\frac{W}{L}\right)}}. \quad (18)$$

With the input voltage high, the output voltage depends upon how much current is steered from the right leg to the left. This is determined by the relative drive strengths of M_1 and M_2 . If M_1 is chosen such that it can readily sink I_0 given the input high gate

voltage, the bias current I_0 will flow through M_1 , resulting in a V_{OL} below the threshold voltage of M_2 .

Similar to the source-coupled pair analysis, the bias current required to drive the assumed load at frequencies up to 1.0 GHz is $I_0 = 1.0\text{mA}$. Assuming a $V_{DSAT} = 0.5\text{V}$, the size of the current source transistor, M_0 , is $\left(\frac{W}{L}\right)_0 = 209$. Since the transistor M_2 is designed to be in cutoff mode when the output is low, and the input transistor M_1 need not remain in saturation, the output voltage can readily sweep through the desired ΔV of 1.0V without violating the circuit's biasing assumptions. This is in marked contrast to the source-coupled pair.

The reduced transistor stack height allows the CSA to operate correctly at lower power supply voltages than are required for the source-coupled pair. Another characteristic of the CSA is that the current source M_0 provides a measure of immunity from power supply variations. As long as the transistor remains in saturation, the bias current is (ideally) independent of the drain-to-source voltage. In real devices, there is some power supply dependence due to channel length modulation, but the sensitivity is quite small. Using the common square law relationship for MOS drain current, the sensitivity of drain current to V_{DS} variation is derived as follows.

$$I_{DS} = \frac{K_P W}{2 L} (V_{GS} - V_T)^2 (1 + \lambda V_{DS})$$

$$S_{V_{DS}}^{I_{DS}} = \frac{V_{DS}}{I_{DS}} \frac{\partial I_{DS}}{\partial V_{DS}} = \frac{V_{DS}}{I_{DS}} \cdot \left[\frac{K_P}{2} \left(\frac{W}{L} \right) (V_{GS} - V_T)^2 \lambda \right]$$

$$S_{V_{DS}}^{I_{DS}} = \frac{V_{DS}}{\frac{K_P W}{2 L} (V_{GS} - V_T)^2 (1 + \lambda V_{DS})} \cdot \frac{K_P}{2} \left(\frac{W}{L} \right) (V_{GS} - V_T)^2 \lambda = \frac{\lambda V_{DS}}{1 + \lambda V_{DS}}$$

Since $\lambda V_{DS} \ll 1$ (typical λ is approximately 0.01), $S_{V_{DS}}^{I_{DS}} \cong \lambda V_{DS}$. Typical variations in V_{DS} are often assumed to be on the order of 10% of V_{DD} . In this case, that would be 0.2 V. With a λ of 0.01, the sensitivity of I_{DS} to V_{DS} is approximately 0.002, or

-54 dB. The dependence of the I_{DS} sensitivity to V_{DS} variation on channel length modulation suggests that the current source transistor M_0 be implemented with long channel devices. Increased channel length will further mitigate the effects of channel length modulation. This is not surprising, as the analysis is identical to that for a simple current mirror.

To further increase the CSA's noise tolerance, more advanced current mirror techniques such as cascode and Wilson current mirror circuits could be used. Such techniques increase the output resistance of the current mirror network, and increase the mirror's tolerance to power supply noise. However, they also require additional voltage headroom.

The current steering amplifier is also tolerant to ground and substrate noise. Since both the input and the output are referred to the same voltage, ground and substrate noise are essentially common mode. The voltage variation affects the gate-to-source voltages of both M_1 and M_2 equally, causing an identical variation in current. This has no effect on the output voltage, and the common mode signal is rejected.

The current steering amplifier is a versatile circuit, which is used in nearly every block of my CMOS PLL. The following sections detail the design and analysis of the individual blocks. The analysis is extended in the case of the current controlled oscillator to include the inherent transistor noise. The derivation in [16] is followed to determine the timing jitter associated with the current steering amplifier.

5.3.1 Charge pump/Loop filter

The detailed circuit discussions begin with the charge pump and loop filter. The charge pump is responsible for accepting signals from the phase-frequency detector and

converting them into current pulses that add or remove charge from the loop filter. The general block diagram for a current charge pump is repeated here as Figure 5.6.

The switches control whether a net charging or discharging current is seen at the output. There are three requirements for an effective charge pump circuit.

1. Equal charge/discharge current regardless of charge pump output voltage.
2. Minimal charge sharing between the output node and the floating nodes created by the open switches.
3. Minimal charge injection from the input signals to the output node.

The reduced operating range of the CGaAs PLL due to charge pump saturation, emphasizes the importance of the first requirement. While that particular implementation was a voltage source charge pump, rather than a current source charge pump, the outcome is similar. As the output of the charge pump changes (in typical designs it tracks the control voltage), either of the current sources may lose the voltage headroom that they require to operate properly. This is most commonly evidenced in the drain to source voltage across a current mirror falling below V_{DSAT} . The behavior of the circuit changes as the device leaves the saturation region.

The commonly employed switch-based circuit is what causes the second requirement. A closed switch allows current to flow either into or out of the charge pump output. When the switch opens, the node between the current source and the switch




Figure 5.6 Generic current source charge pump block diagram.

becomes a floating node. The node voltage is typically pulled near the appropriate rail voltage (V_{DD} for the charging switch, ground for the discharging switch). As Figure 5.7 demonstrates, when the switch closes, there is now a connection between two nodes with different voltages. This causes charge sharing between the two nodes which produces a perturbation on the charge pump output voltage.

The standard charge sharing analysis estimates this perturbation as follows. Modeling the floating node and loop filter simply as two capacitors C_{FN} and C_1 , respectively, a switch closing between them causes the voltage to equalize across them. The total capacitance becomes $C_F = C_{FN} + C_1$. The final charge is the sum of the initial charges, $C_{FN}V_{FN} + C_1V_1$, since charge must be conserved. Therefore the final voltage $V_F = \frac{Q_F}{C_F} = \frac{C_{FN}V_{FN} + C_1V_1}{C_{FN} + C_1}$. However, the loop filter capacitor $C_1 \gg C_{FN}$ so $V_F \approx \frac{C_{FN}}{C_1}V_{FN} + V_1$. This results in a $\Delta V = \frac{C_{FN}}{C_1}V_{FN}$. This perturbation is quite small, as C_1 can easily be three orders of magnitude greater than C_{FN} . Yet it is an additive effect that occurs every cycle, so it must be considered and minimized. As stated previously, this is another reason for choosing a large loop filter capacitance.

Charge injection occurs due to coupling from the input signals to the output node through the gate to drain capacitance of the switch transistors. This effect is made more




Figure 5.7 Illustration of charge sharing within the charge pump.

pronounced by large swings on the inputs, and large switching transistors (which result in a larger gate to drain capacitance). From the illustration in Figure 5.8, it is apparent that the situation essentially imitates a capacitive divider. Given ΔV_{in} , the perturbation on the filter voltage V_1 is approximately $\Delta V_1 = \frac{C_{gd}}{C_{gd} + C_1} \Delta V_{in}$. Since $C_1 \gg C_{gd}$ this relation reduces to $\Delta V_1 \approx \frac{C_{gd}}{C_1} \Delta V_{in}$. Again, this is not a very significant amount, but it has an additive effect that occurs every time the charge pump input signal transitions.

Minimizing this effect implies both maximizing the loop filter capacitance in relation to the coupling capacitance, and minimizing the voltage swing at the charge pump input.

As the following discussion will reveal, the current steering amplifier lends itself well to this application. While typical designs utilize source-coupled pairs (which suffer from charge injection), or pass-gate style switches (which suffer from both charge sharing and charge injection) the CSA implementation minimizes charge injection and completely eliminates charge sharing. Furthermore, the use of an active integrator ensures equal charging and discharging currents through the full range of output control voltages, as will be shown shortly.

The concept behind the CSA charge pump is that the current flowing through the load device is readily mirrored. Given this, it is possible to implement a circuit that

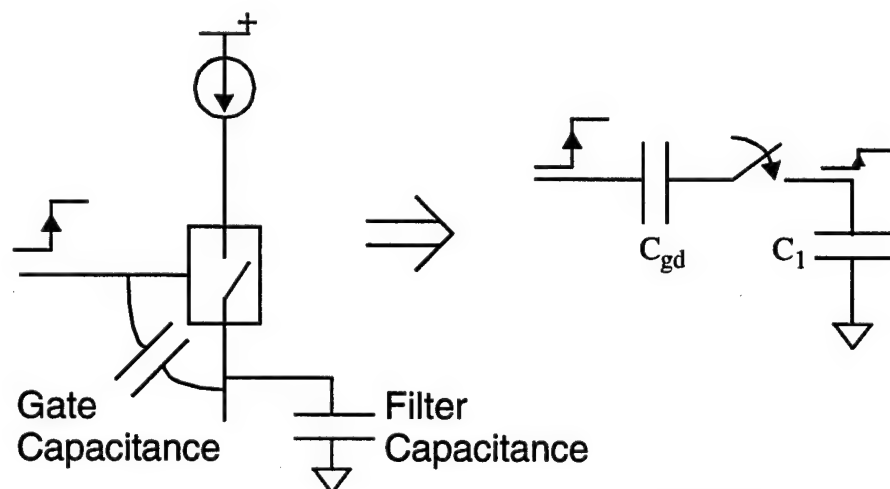


Figure 5.8 Illustration of charge injection in a charge pump.

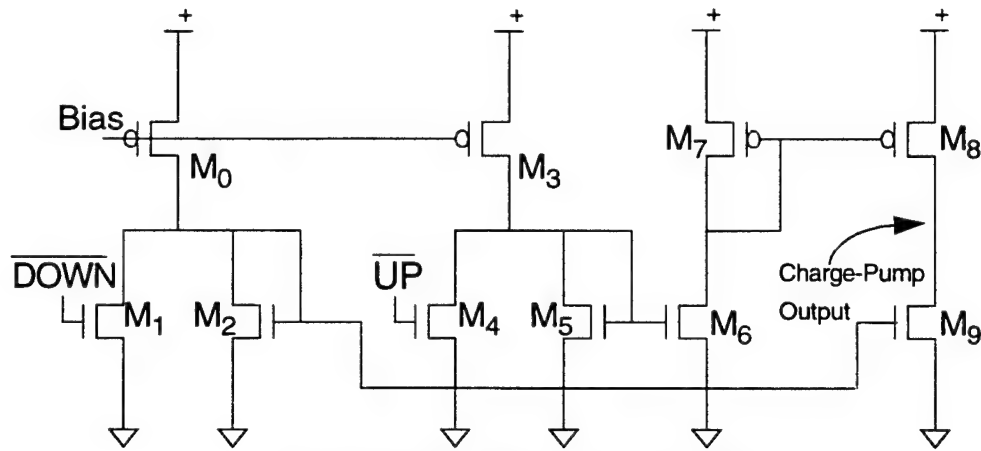


Figure 5.9 CSA charge pump schematic.

simply mirrors CSA load currents to an output stage that produces charging and discharging currents. Figure 5.9 shows the complete charge pump circuit.

Operation of the charge pump, illustrated in Figure 5.10, is summarized as follows. Biased at a current $I_0 = 2I_{CP}$, the driver and load transistors are sized such that an input high voltage draws approximately I_{CP} through each leg. The load currents of both CSA's are mirrored to the output stage of M_8 and M_9 . With equal currents mirrored to M_8 and M_9 , the net current at the output is zero. When the DOWN input pulses low, the input device, M_1 , cuts off, steering $2I_{CP}$ through the load device M_2 . This is mirrored to M_8 . The current through M_9 remains I_{CP} , but the current through M_8 is biased to $2I_{CP}$. This produces a net discharge current of I_{CP} . Conversely, a downward pulse on the UP signal produces a net charging current of I_{CP} . Since this implementation does not employ any




Figure 5.10 Graphical illustration of CSA charge pump operation.

switches, there are no floating nodes, and therefore no charge sharing. Additionally, with a specified I_{CP} , the sizes of the current mirror transistors M_2 and M_7 control the voltage swing at the inputs to the output transistors M_8 and M_9 . The voltage swing representing the switch between I_{CP} and $2I_{CP}$ is on the order of 0.1V. This results in a 95% reduction in charge injection over implementations that use rail-to-rail switching signals.

The first requirement for an effective charge pump is met through the use of an active integrator. The block diagram in Figure 5.11 shows the simple topology. The integrator produces an output voltage given by the following expression.

$$V_{out} = - \int I_{CP} dt + V_{C1} \quad (19)$$

The operational amplifier is a simple source-coupled pair. The primary reason to use such an implementation is that the input node, which is essentially the output node of the charge pump, is held at V_{REF} by the action of the negative feedback. This occurs regardless of the integrator's output voltage. Thus, the output voltage of the charge pump is invariant with respect to the VCO control voltage, eliminating the problem of reduced current drive at the control voltage extremes.

To further aid current matching, the reference voltage is generated via a replica circuit to match the drain voltage of M_7 . This ensures that M_7 and M_8 are operating with




Figure 5.11 Active loop filter implementation.

nearly identical drain to source voltages, reducing the effect of channel length modulation on the current mirror matching.

Sizing of the charge pump begins with the current source transistor, M_0 . This device is sized such that it can remain safely in saturation during normal operation. With $I_{CP} = 15 \mu A$, the device M_0 must be sized for $I_0 = 30 \mu A$.

$$I_0 = 2I_{CP} = 30\mu A = \frac{1}{2}K_P\left(\frac{W}{L}\right)_0 (V_{SG0} - |V_{tp}|)^2, V_{DSAT} = 0.3 \Rightarrow \left(\frac{W}{L}\right)_0 = 15.5$$

Next, the transistors M_1 and M_2 are equally sized, such that a nominal voltage level at the output is achieved with the same voltage at the input, and I_{CP} flowing through each device. This voltage is chosen to be $V_{OUT} = 1.2 V$. Since both the input and output will be biased at 1.2 V, M_1 and M_2 are saturated.

$$I_1 = I_2 = I_{CP} = 15\mu A = \frac{K_P W}{2 L} (V_{GS} - V_T)^2$$

$$15\mu A = \frac{1}{2}(1.78 \times 10^{-4})\left(\frac{W}{L}\right)(1.2 - 0.67)^2 \Rightarrow \left(\frac{W}{L}\right)_1 = \left(\frac{W}{L}\right)_2 = 0.6$$

Identical sizing applies to the CSA given by $M_3 - M_5$. The mirror devices M_6 and M_7 are chosen to exactly mirror the current through M_5 to M_8 . Similarly, the discharge current is mirrored from M_2 to M_9 .

The plot in Figure 5.12 shows a Saber simulation of the charge pump operation. The plot shows the action of the UP and DOWN inputs, and the resulting net output current. The figure depicts circuit operation when the UP input is active (net charging current), and the DOWN input is active (net discharging current).

Another useful simulation is to run the charge pump as if the PLL is in the locked condition with short, simultaneous pulses on the UP and DOWN inputs. The result of this simulation is presented in Figure 5.13 (on p. 82). The average value of the charge/discharge current is an important parameter. The average current represents the error in




Figure 5.12 Example of charge pump operation.

current matching through the charge pump. This charge pump implementation achieves 0.0478 % error.

So far, the discussion has centered upon the design of the primary charge pump. The auxiliary charge pump is identical, but the ability to change I_{CP} is needed to provide the means of varying A_f . The charge pump implementation described above lends itself very well to a simple solution to this problem.

The current-steering amplifiers were sized such that an input voltage of 1.2 V would split the source current evenly between M_1 and M_2 . If the input voltage were increased to 1.4V, however, this would steer more current away from the load transistor when the input signal is inactive. Note that this has no effect on the output current until one of the inputs transitions low. At this point, all of the current flows through the appropriate load transistor and is mirrored to the output. However, in this case, the current




Figure 5.13 Charge pump output current in the phase-locked state.

mirrored to the inactive output device is less than I_{CP} , so the net output current pulse is greater than I_{CP} . Conversely, if the input voltage is reduced to 1.0V, the net output current pulse will be smaller than I_{CP} . This is a simple means of controlling the feed forward current gain, A_f . Furthermore, the output voltage level of the phase frequency detector is readily controlled, once again due to the flexibility of the current-steering amplifier. This is discussed in the next section.

5.3.2 Phase-frequency Detector

Operation of the phase frequency detector has been well covered in previous sections. The simple implementation of two D-type flip-flops and an AND gate is improved upon in [3]. Figure 5.14 shows the PFD logic diagram used in this design. The DFF's in the conventional implementation have their data inputs tied to a logical one, as illustrated in Figure 2.15 (on p. 17). Minimizing the logic depth within the DFFs, given the constant input, improves the logic delay through the circuit. It is desirable for a phase-frequency detector to output short, simultaneous pulses on both outputs when the input signals are in phase. The minimum pulse width depends upon the delay through the PFD logic. The absence of a dead zone also characterizes this PFD implementation.

To further minimize the phase jitter contribution of the phase-frequency detector, the current steering amplifier was adapted to implement logic functionality by replacing the single input transistor with an NMOS network. Sizing conventions for the n-transistors are followed to retain the appropriate overall drive strength. For example, if a regular CSA has a driver size of $\left(\frac{W}{L}\right)$, a NAND3 implementation would require three




Figure 5.14 Phase-frequency detector block diagram.




Figure 5.15 AOI21 CSA logic gate schematic.

series input devices of size $3\left(\frac{W}{L}\right)$. Figure 5.15 is an example of a current steering logic gate.

To make these gates compatible with the charge pump, it is necessary to have a variable output high voltage. In the basic CSA, the output high voltage is set by choosing an appropriate size for the diode-connected load transistor. This implementation obviously does not lend itself well to variation.

The diode connected load transistor can be replaced by a resistor without changing the circuit's functionality. If it is replaced with a variable resistor, the output high voltage becomes variable as well.

Figure 5.16 illustrates a CSA gate with a voltage variable resistor load. This resistor implementation is identical to that reported in [15]. The voltage variable resistor

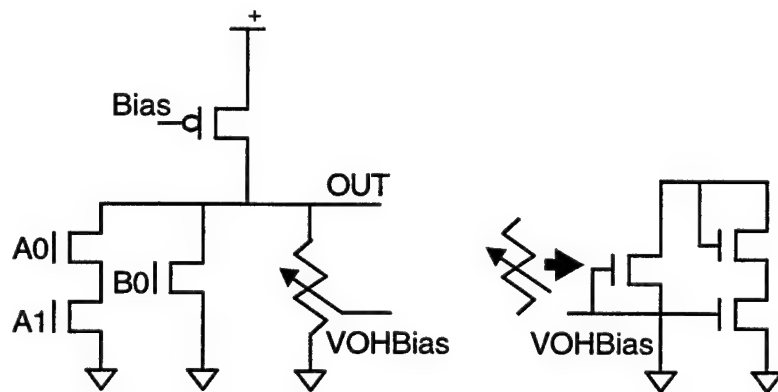


Figure 5.16 CSA logic gate with V_{OH} control.




Figure 5.17 Regulation of V_{OH} using replica feedback biasing.

provides a more linear voltage-to-current characteristic than a single transistor. The effective resistance changes with a change on the input bias voltage.

To provide a robust and stable resistor bias that is relatively tolerant to process parameter variation a negative feedback loop is used. This negative feedback utilizes a CSA replica with its input tied low, as illustrated in Figure 5.17. The low input causes the CSA to output a high voltage. The negative feedback drives the output high voltage to equal the input reference voltage through variation of the resistor bias voltage. The reference voltage originates in the bias generator block, as will be discussed in Section 5.3.6.

The Saber output in Figure 5.18 shows the phase-frequency detector operation at reference voltages of 1.0, 1.2, and 1.4 volts. Note that the output high voltages closely match the desired level. These gates are used throughout the phase-frequency detector, but the output gates that produce the UP/DOWN signals are replicated. These duplicate outputs are biased separately since only the outputs to the auxiliary charge pump require programmability of the output high voltage.

The specification of device sizes for the various components of these CSA logic gates follows a process similar to that discussed for the generic CSA (which is a CSA inverter). The voltage-controlled resistor sizes are set such that the voltage-to-current




Figure 5.18 Example of PFD operation with V_{OH} control.

characteristic is sufficiently linear over the desired voltage range (1.0 V to 1.4 V). Figure 5.19 illustrates the resulting voltage-to-current characteristic for the voltage-controlled




Figure 5.19 Voltage to current characteristic for the voltage-controlled resistor.

resistor. These sizes were determined through simulation to be those depicted in the circuit schematic of Figure 5.20.

The current source transistor size follows the same analysis as the generic CSA. The target bias current is $I_0 = 100 \mu\text{A}$, and the device must remain in saturation with a $V_{DS} = 0.4 \text{ V}$. This arises from the fact that the maximum output voltage is 1.4 V , with a 0.2 V margin designed for safety. The input devices sink the $100 \mu\text{A}$ bias current and produce an output low voltage when active. In this state the input transistor is in the linear mode. Assuming an output low voltage of 0.3 V , the following relation determines the required size. If more than a simple inverting logic function is needed, this size is scaled accordingly.

$$I_D = 100 \mu\text{A} = K_P \left(\frac{W}{L} \right) \left[(V_{GS} - V_T) V_{DS} - \frac{1}{2} V_{DS}^2 \right]$$

$$= (1.78 \times 10^{-4}) \left(\frac{W}{L} \right) \left[(1.2 - 0.67)(0.3) - \frac{1}{2} (0.3)^2 \right] \Rightarrow \left(\frac{W}{L} \right) = 5.0$$

Implementing the PFD with the sizes shown in Figure 5.20, the minimum pulse width can be plotted over several values of bias current. This provides a view of the trade-




Figure 5.20 CSA logic gate sizes used in the PFD.




Figure 5.21 Minimum PFD pulse width versus PFD bias current.

off between performance and power dissipation. The resulting plot of Figure 5.21 shows the reduced slope of the curve at bias currents above $80\ \mu\text{A}$. Since a minimum pulse width of $700\ \text{ps}$ is sufficiently small, the bias current level was set at $80\ \mu\text{A}$. Device sizes were left largely the same and the I-V characteristics of the voltage-controlled resistor were verified at this current level.

A simulation measuring the net output pulse width for varying input phase errors determines the magnitude of the dead zone in a phase-frequency detector design. In the simulation, the designer applies clock signal inputs that are separated by small phase errors. The dead zone, if one exists, will manifest around the point of zero input phase error. By plotting the difference in output pulse width versus input phase error, any regions with low, or even zero, gain are revealed. Figure 5.22 shows that the CSA phase-frequency detector has no dead zone in its transfer characteristic.




Figure 5.22 Net pulse width versus input phase error

5.3.3 Current-Controlled Oscillator

The most sensitive circuit in the PLL, the current-controlled oscillator requires detailed analysis to minimize its contribution to overall phase jitter. As per the previous discussions, the ICO is implemented as a 3-stage ring oscillator of CSA stages. Reasons for this choice include the CSA's tolerance to power supply variation, its ability to operate at low power supply voltages and high frequencies, and its transistor efficient implementation.

Analysis of the ICO begins with the derivation of the current steering amplifier's inherent timing jitter. The circuit diagram depicted in Figure 5.23 serves as the subject for

this analysis. The steps follow those presented in [16], as demonstrated in Section 3.2.

They include:

- Step 1. Determine the stage delay, t_d ;
- Step 2. Find the equivalent noise generators;
- Step 3. Relate voltage noise to phase jitter with the first crossing approximation;
- Step 4. Determine the interstage interaction;
- Step 5. Extend the phase jitter of a single stage to that of a ring oscillator.

Step #1: As the schematic in Figure 5.23 illustrates, the capacitance at the output of the CSA stage is modeled by the lumped capacitor, C_L . Assuming that the following stage begins switching when the output passes the midpoint of its swing, equation (18) approximates the stage delay.

$$t_d \cong \frac{1}{2} \Delta V \frac{C_L}{I_o} \quad (20)$$

The voltage swing, ΔV , is determined from the current equations. With an input high voltage, the resultant output voltage is low, resulting in the following current equation for the driver transistor, M_1 .

$$I_o = \mu C_{ox} \left(\frac{W}{L} \right)_1 \left[(V_{OH} - V_T) V_{OL} - \frac{1}{2} V_{OL}^2 \right]$$

$$V_{OL}^2 - 2(V_{OH} - V_T) V_{OL} + \frac{2I_o}{\mu C_{ox} \left(\frac{W}{L} \right)_1} = 0$$

This equation can be solved for V_{OL} using the quadratic equation. The difference

between V_{OH} (given in equation (18)), and the V_{OL} , is the voltage swing.

$$V_{OL} = \frac{V_T - \sqrt{4(V_{OH} - V_T)^2 - 4 \frac{2I_0}{\mu C_{ox} \left(\frac{W}{L}\right)_1}}}{2}$$

$$\Delta V = V_{OH} - V_{OL} = V_T + \sqrt{(V_{OH} - V_T)^2 - \frac{2I_0}{\mu C_{ox} \left(\frac{W}{L}\right)_1}}$$

$$\Delta V = V_T + \sqrt{\frac{2I_0}{\mu C_{ox} \left(\frac{W}{L}\right)_2} - \frac{2I_0}{\mu C_{ox} \left(\frac{W}{L}\right)_1}} = V_T + \sqrt{\frac{\left(\frac{W}{L}\right)_1 - \left(\frac{W}{L}\right)_2}{\left(\frac{W}{L}\right)_1 \left(\frac{W}{L}\right)_2} \frac{2I_0}{\mu C_{ox}}} \quad (21)$$

Step 2: The four current noise generators depicted in Figure 5.23 represent the relevant noise sources for this problem. They sum readily at the output to result in the noise current given by equation (21).

$$\overline{i_n^2} = \overline{i_{n1}^2} + \overline{i_{n2}^2} + \overline{i_{n3}^2} + \left(\frac{W_3}{W_4}\right)^2 \overline{i_{n4}^2} \quad (22)$$

Where, $\overline{i_{nx}^2} = 4kT \frac{2}{3} g_m \Delta f$ for saturated devices, and $\overline{i_{nx}^2} = 4kT \frac{1}{R_{eff}} \Delta f$ for




Figure 5.23 CSA VCO delay stage with relevant noise current sources.

devices operating in the linear region. The factor $\left(\frac{W_3}{W_4}\right)$ represents the amount of current scaling from the bias source to the CSA circuit.

The overall voltage noise is found by determining the effective resistance seen at the output node, multiplying by the expression for the noise current source, and integrating over the noise bandwidth. Since the circuit is most susceptible to phase jitter during the transitions of the output, the effective resistance and voltage noise source are calculated at the point where the output crosses the midpoint voltage. Thus, the point where $V_{in} = V_{out}$ must be determined.

If $V_{in} = V_{out}$, then $V_{GS1} = V_{DS1}$ and M_1 is saturated. Since, M_2 is diode connected, it too is saturated.

$$g_{m1} = \mu C_{ox} \left(\frac{W}{L}\right)_1 (V_m - V_T) \quad g_{m2} = \mu C_{ox} \left(\frac{W}{L}\right)_2 (V_m - V_T) \quad (23)$$

$$\begin{aligned} \therefore V_{out} &= I_0 \left(\frac{1}{g_{m1}} \parallel \frac{1}{g_{m2}} \right) = V_m \\ &= I_0 \frac{1}{\mu C_{ox} (V_m - V_T) \left[\left(\frac{W}{L}\right)_1 + \left(\frac{W}{L}\right)_2 \right]} \\ V_m^2 - V_T V_m - \frac{I_0}{\mu C_{ox} \left[\left(\frac{W}{L}\right)_1 + \left(\frac{W}{L}\right)_2 \right]} &= 0 \\ \Rightarrow V_m &= \frac{1}{2} V_T + \sqrt{\frac{1}{4} V_T^2 + \frac{I_0}{\mu C_{ox} \left[\left(\frac{W}{L}\right)_1 + \left(\frac{W}{L}\right)_2 \right]}} \end{aligned} \quad (24)$$

At this midpoint, the output resistance is given by (25).

$$R_{out} = r_{o1} \parallel \frac{1}{g_{m2}} \quad r_{o3} \equiv \frac{1}{g_{m2}} \quad (25)$$

The noise bandwidth is estimated by the low pass filter bandwidth formed by the output resistance and the output load capacitance. Using this concept, the noise bandwidth is

given as follows.

$$BW = \frac{\pi}{2} \cdot \frac{1}{2\pi} \cdot \frac{1}{R_{out} C_L} = \frac{1}{4R_{out} C_L} \quad (26)$$

At the midpoint voltage, all the transistors are operating in saturation, so the current noise generator is

$$\overline{i_n^2} = 4kT\frac{2}{3}g_{m1}\Delta f + 4kT\frac{2}{3}g_{m2}\Delta f + 4kT\frac{2}{3}g_{m3}\Delta f + \left(\frac{W_3}{W_4}\right)^2 4kT\frac{2}{3}g_{m4}\Delta f \quad (27)$$

$$\frac{\overline{v_n^2}}{\Delta f} = R_{out}^2 \overline{i_n^2} = \left(\frac{1}{g_{m2}}\right)^2 4kT\frac{2}{3} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4}\right)^2 g_{m4}\right) \quad (28)$$

Integrating this noise power over the noise bandwidth yields the total voltage noise power. Taking the square root results in the voltage noise.

$$\begin{aligned} \overline{v_n^2} &= \frac{1}{4} \left(\frac{g_{m2}}{C_L}\right) \cdot \left(\frac{1}{g_{m2}}\right)^2 4kT\frac{2}{3} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4}\right)^2 g_{m4}\right) \\ v_n &= \sqrt{\frac{2kT}{C_L}} \sqrt{\frac{1}{3} \frac{1}{g_{m2}} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4}\right)^2 g_{m4}\right)} \end{aligned} \quad (29)$$

Step 3: The first crossing approximation estimates the timing variance as $\overline{\Delta t^2} \equiv \overline{\Delta v_n^2} \left(\frac{C_L}{I_0}\right)^2$. Normalizing the timing jitter to the stage delay, t_d , and applying the

results derived above for the voltage noise, gives the following expression for the timing jitter.

$$\frac{\Delta\tau_{1rms}}{t_d} \cong \frac{v_n \cdot \frac{C_L}{I_0}}{\frac{1}{2}\Delta V \cdot \frac{C_L}{I_0}} = \frac{v_n}{\frac{1}{2}\Delta V} \quad (30)$$

$$\therefore \frac{\Delta\tau_{1rms}}{t_d} \cong \frac{\sqrt{\frac{2kT}{C_L}} 2\xi}{\Delta V} \quad (31)$$

$$\text{where, } \xi = \sqrt{\frac{1}{3} \left[\frac{1}{g_{m2}} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4} \right)^2 g_{m4} \right) \right]} \quad (32)$$

Step 4: The interstage interaction is a second order effect that must be accounted for in order to preserve the accuracy of the analysis. Figure 5.24 shows both the block




Figure 5.24 Interstage interaction.

diagram and small-signal circuit schematic for the interstage scenario. From the circuit diagram, the interstage interaction is readily derived.

$$v_{n1} = i_n R_{out} \quad \text{with } G_m = g_{m1}, \quad v_{n2} = g_{m1} R_{out}^2 i_n$$

Converting this result into a spectral noise density allows the noise power generated by the previous stage to be simply added to that generated in the main stage.

$$\overline{v_{n2}^2} = [g_{m1} R_{out}]^2 \overline{i_n^2}$$

$$\therefore \overline{v_n^2} = [R_{out}^2 + (g_{m1} R_{out}^2)^2] \overline{i_n^2} = (1 + (g_{m1} R_{out})^2) R_{out}^2 \overline{i_n^2} = (1 + a_v^2) R_{out}^2 \overline{i_n^2}$$

As before, $\overline{i_n^2} = 4kT \frac{2}{3} g_{m1} \Delta f + 4kT \frac{2}{3} g_{m2} \Delta f + 4kT \frac{2}{3} g_{m3} \Delta f + \left(\frac{W_3}{W_4}\right)^2 4kT \frac{2}{3} g_{m4} \Delta f$, and the noise bandwidth, $BW = \frac{1}{4} \frac{1}{R_{out} C_L}$. Combining these yields the following noise voltage with interstage interaction accounted for.

$$v_n = \sqrt{1 + a_v^2} \sqrt{\frac{2kT}{C_L}} \sqrt{\frac{1}{3} \frac{1}{g_{m2}} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4}\right)^2 g_{m4} \right)} \quad (33)$$

This results in the following expression for the timing jitter normalized to the stage delay.

$$\begin{aligned} \frac{\Delta\tau_{1rms}}{t_d} &\equiv \frac{v_n \cdot \frac{C_L}{I_0}}{\frac{1}{2} \Delta V \cdot \frac{C_L}{I_0}} = \frac{v_n}{\frac{1}{2} \Delta V} \\ \therefore \frac{\Delta\tau_{1rms}}{t_d} &\equiv \sqrt{1 + a_v^2} \sqrt{\frac{2kT}{C_L}} \frac{2\xi}{\Delta V} \end{aligned} \quad (34)$$

Where the factor, ξ , is that of equation (31). Thus the effect of the interstage interaction is to contribute a component of timing jitter that is proportional to the voltage gain through

a single stage.

Step 5: With the timing jitter for a single stage, normalized to the stage delay, given above, the final step is to extend this value to the timing jitter for a full ring oscillator composed of N such stages.

$$\overline{\Delta\tau_N^2} = \overline{\Delta\tau_1^2} \times \frac{T_0}{t_d} = (1 + a_v^2) \frac{2kT}{C_L} \xi^2 \frac{1}{\left(\frac{1}{2}\Delta V\right)^2} T_0 \frac{C_L}{I_0} \cdot \frac{1}{2}\Delta V$$

$$\overline{\Delta\tau_N^2} = \frac{4kT}{I_0} (1 + a_v^2) \xi^2 \frac{T_0}{\Delta V} \quad (35)$$

Similar to the result for the source-coupled pair presented in Section 3.2 from [16], the ring oscillator phase jitter is inversely proportional to the bias current, establishing a phase jitter/power dissipation trade-off. In this case, however, the bias current is determined by the frequency specifications. The noise factor ξ , is minimized if $g_{m2} \gg g_{m3}$ and $g_{m2} \gg g_{m1}$. It is also desirable to maximize the current-controlled oscillator (ICO) output voltage swing.

Since the worst case biasing condition occurs at the maximum current end of the ICO operating range, the current source transistor is sized so that it remains in saturation at this current (1 mA). The $\left(\frac{W}{L}\right)$ ratio of the current source transistor is determined assuming a $V_{DSAT} = 0.5$ V.

$$I_0 = 1\text{mA} = \frac{1}{2} \left(\frac{W}{L}\right) K_P (V_{DSAT})^2 \Rightarrow \left(\frac{W}{L}\right)_0 = 186$$

Similarly, M_2 is sized to sink I_0 , providing a nominal output high voltage of 1.2 V. This calculation results in a $\left(\frac{W}{L}\right)_2 = 40$. The remaining device size is set by the desired




Figure 5.25 Ring oscillator schematic.

output low voltage. The choice of $V_{OL} = 0.2$ V sets its size through the linear transition equation at $\left(\frac{W}{L}\right)_1 = 65$.

The specifications related to bias current range assumed a specific load capacitance. Saber simulations predicted that the ICO of Figure 5.25 would achieve the desired frequency range.

Since the input device is never used in saturation, short channel effects are of no concern and a minimum channel length device is used. This helps reduce the load




Figure 5.26 ICO frequency versus bias current.

capacitance. 1 μm channel lengths are used for the other two transistors to reduce the non-idealities. The plot in Figure 5.26 shows the bias current to frequency characteristic of the oscillator using the initial sizes.

Repeating the current simulation vs. frequency over a range of power supply voltages demonstrates the current steering amplifier's tolerance to power supply variation. The curves in Figure 5.27 represent the bias current to frequency characteristic for power supply voltages of 1.8, 2.0, and 2.2 volts. Note that throughout the operating frequency range, the oscillator frequency is largely independent of power supply voltage.

While the oscillator is the most sensitive circuit, it is also the most computationally intensive in terms of simulation time. PLLs are difficult to simulate because they typically require picosecond time-steps over microsecond time-frames. The small time-step is due, primarily, to the high frequency of the voltage-controlled oscillator. This severely limits the range of full PLL simulations. A piecewise linear behavioral model of the oscillator is therefore very useful. The plot in Figure 5.28 shows the frequency-to-current




Figure 5.27 Frequency-to-current characteristic of the ICO over V_{DD} .




Figure 5.28 Piecewise linear ICO model.

characteristic of the piecewise linear oscillator model. The use of a piecewise linear model is even more effective when it includes the frequency divider as a $1/N$ scaling factor of the frequency. Using this model reduces the run-time of full system simulations by an order of magnitude.

5.3.4 V-I Converter

The voltage-to-current converter is arguably the second most sensitive circuit in the PLL. The output bias that sets the oscillator frequency must remain stable with respect to power supply voltage. Furthermore, since the oscillator current is set by a PMOS

device (whose gate voltage is referenced to V_{DD}), the V-I converter's output voltage must also be referenced to V_{DD} .

Another desirable characteristic for a V-I converter is that it complement the frequency-to-current transfer function of the oscillator, such that the combination results in a linear voltage-to-frequency relationship. To estimate the frequency-to-current relationship for a ring oscillator with N stages, it is convenient to start with the delay through a single stage, $\Delta t \equiv \frac{C_L}{I_0} \Delta V$. This implies that $f = \frac{1}{2N\Delta t} = \frac{I_0}{2NC_L\Delta V}$. The voltage swing, to a first order, is given by the CSA V_{OH} , which is proportional to the square root of the bias current, I_0 . This results in a frequency which is also proportional to the square root of the bias current. To produce an overall linear relationship, the output current of the voltage-to-current converter should be proportional to the square of the input voltage. A MOS transistor, biased in saturation, provides this relationship.

One implementation of a V-I converter that meets these guidelines is shown in Figure 5.29. The circuit is an active current mirror whose current is set by an input voltage.

The opamp is configured in a negative feedback loop. The feedback equalizes the drain voltages of the two NMOS devices and produces the necessary bias voltage to mirror the current through M_1 in the two PMOS devices. Choosing the size of M_2 such that the required V_{GS} will keep M_1 in saturation, assures the output current's dependence on the square of the input voltage. The transistor M_1 is sized to produce the high range of




Figure 5.29 Active current mirror V-I converter.

oscillator bias current at the high end of the input control voltage. To minimize power dissipation, the V-I converter generates a smaller current than what is needed in the oscillator stages. The V-I converter is a current mirror that outputs the bias voltage for the ICO current source; simple current mirror scaling provides the correct range of currents at the ICO.

The output of the auxiliary charge pump is fed directly into the drain of M_1 , allowing the auxiliary charge pump to perturb the bias current, and complete the feed forward loop. The forward gain, A_f , is set by the variation of I_{CP} from the nominal value, and the current scaling factor from the V-I converter to the oscillator. This scaling factor is chosen to be 6.67, so that A_f ranges in value from 2 to 9.

Calculating initial sizes for the V-I converter and simulating its DC behavior produced good results. The plot in Figure 5.30 demonstrates the square-law relationship




Figure 5.30 Active current mirror V-I converter DC transfer characteristic.

between the V-I converter and ICO bias currents and the input control voltage. The circuit's transient behavior, however, exposes a problem with the implementation.

If a step function is applied to the power supply voltage, the active current mirror takes a finite period of time to correct the current bias. This occurs quickly, but not before a number of oscillator cycles pass. The resulting frequency perturbation, as illustrated in Figure 5.31, produces significant phase jitter. Even with proper compensation, the problem persists.

The lower current level of the V-I converter allows it to be implemented as a differential pair. The differential design achieves power supply independence without the problems associated with the active current mirror. The circuit diagram in Figure 5.32 illustrates the differential V-I converter.

The implementation takes advantage of the current-splitting properties of the source-coupled pair. The tail current, set by the current source transistor, M_5 , is divided




Figure 5.31 Active current mirror V-I converter output voltage instability.




Figure 5.33 Source-coupled pair V-I converter transfer characteristic.

provide a very stable output bias voltage in response to a step in power supply voltage.

The trace in Figure 5.34 demonstrates this stability.




Figure 5.34 Power supply step response of the differential V-I converter.

5.3.5 Frequency Divider

The frequency divider is implemented as a series of toggle flip-flops. Each successive stage provides another division factor of 2. The counter outputs include a $\div 2$ clock, a $\div 4$ clock, and a $\div N$ clock. The $\div N$ clock is the result of a multiplexor select between the divide ratios of 2, 4, 8, 16, and 32. It is the $\div N$ output that feeds back to the phase-frequency detector.

The flip-flops are implemented as sense-amp flip-flops, as reported in [32]. Figure 5.35 illustrates the flip-flop circuit diagram. Two types of flip-flops are used, differing in the implementation of the NAND gates. As the logic diagram of the frequency divider in Figure 5.36 shows, three of the flip-flops are clocked by the high frequency oscillator output. These flip-flops require a faster implementation, so their NAND gates utilize CSA logic. The remaining registers employ simple, complementary NAND gates.

The $\div 2$ and $\div 4$ output clock signals are synchronized with the oscillator output. Since the PLL clock generator was designed to provide the $\div 1$, $\div 2$, and $\div 4$ clock signals to different chips in a multi-chip system, the first order synchronization of the three outputs is important.




Figure 5.35 Sense-amp D-type flip-flop schematic.

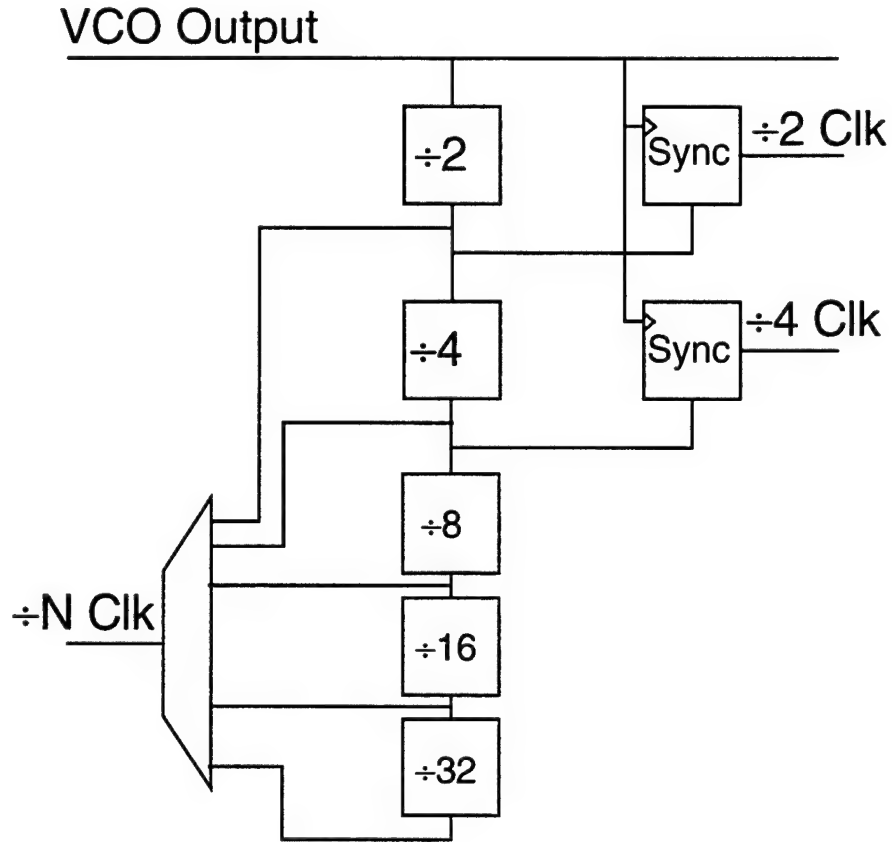


Figure 5.36 Frequency divider block diagram.

5.3.6 Bias Generator

Each of the PLL components discussed requires one or more bias voltages for correct operation. The bias generator provides all of these required voltages. The list of required bias voltages is as follows:

- NMOS input opamp current bias (30 μ A);
- PMOS input opamp current bias (30 μ A);
- Charge pump current bias (30 μ A);
- Phase-frequency detector current bias (80 μ A);
- V-I converter static bias (1.4 V);
- Auxiliary charge pump V_{OH} reference (1.0 V - 1.4 V);
- Main charge pump V_{OH} reference (1.2 V);




Figure 5.37 ΔV_{BE} bias generator concept illustration.

The bias generator requires careful design as errors in these bias outputs are likely to degrade the overall PLL performance, if not cause the PLL to fail altogether. As with the other circuit blocks, tolerance to power supply variation is of primary importance. It should be noted, however, that more power supply dependence is tolerable in the bias generator. This is in contrast to such blocks as the oscillator and V-I converter, where changes in the power supply directly affect the frequency of the oscillator.

The heart of the circuit is a kT/q generator. The circuit in Figure 5.37 shows the concept behind a kT/q generator. Also known as a ΔV_{BE} generator, the circuit uses ratioed bipolar junction transistors to establish a known voltage across a resistor.

If the voltages at nodes 1 and 2 are held equal, so that the currents through the legs are equal, then the following relationship exists.

$$V_1 = V_2 \Rightarrow V_{BE1} = IR + V_{BE2} \Rightarrow V_T \ln\left(\frac{I}{I_s}\right) = IR + V_T \ln\left(\frac{I}{NI_s}\right) \Rightarrow$$

$$IR = V_T \ln\left(\frac{I}{I_s}\right) - V_T \ln\left(\frac{I}{NI_s}\right) = V_T \ln(N) \Rightarrow I = \frac{V_T \ln(N)}{R} \quad (36)$$

As the final relation indicates, and the name implies, the current generated by this circuit is a function of the two bipolar devices. Given that the current in the two devices is equal, the larger BJT will require less V_{BE} to realize that current. The difference,

$\Delta V_{BE} = V_{BE1} - V_{BE2}$, falls across the resistor, providing a means to reliably and accurately set the current value, within the accuracy of the resistor fabrication process.

The other half of the kT/q generator is the current mirror circuitry which equalizes the current in the two legs. In this implementation, an active current mirror is used. The circuit in Figure 5.38 equalizes the voltages at nodes 1 and 2 through negative feedback, while the PMOS current mirror equalizes the currents through the two legs of the circuit.

The PNP BJTs are implemented with P+ islands in an N-well, while the resistor is a patterned polysilicon resistor. Since both the BJT size ratio and the resistor value set the current, the PMOS transistors can be sized appropriately to mirror the current to the destination PLL components.

To minimize the number of current mirroring branches in the bias generator, the design uses two distinct kT/q generators. The circuit diagrams in Figure 5.39 and Figure 5.40 represent the entire bias generator schematic.

The left half of the circuit is responsible for generating the charge pump current bias, the opamp current biases, and the V-I converter bias voltage. As such the ratio of BJT sizes, $N=10$ and $R=1.9K\Omega$ results in a current of $I = 31.3 \mu A$. This is close enough to




Figure 5.38 Complete kT/q bias generator schematic.

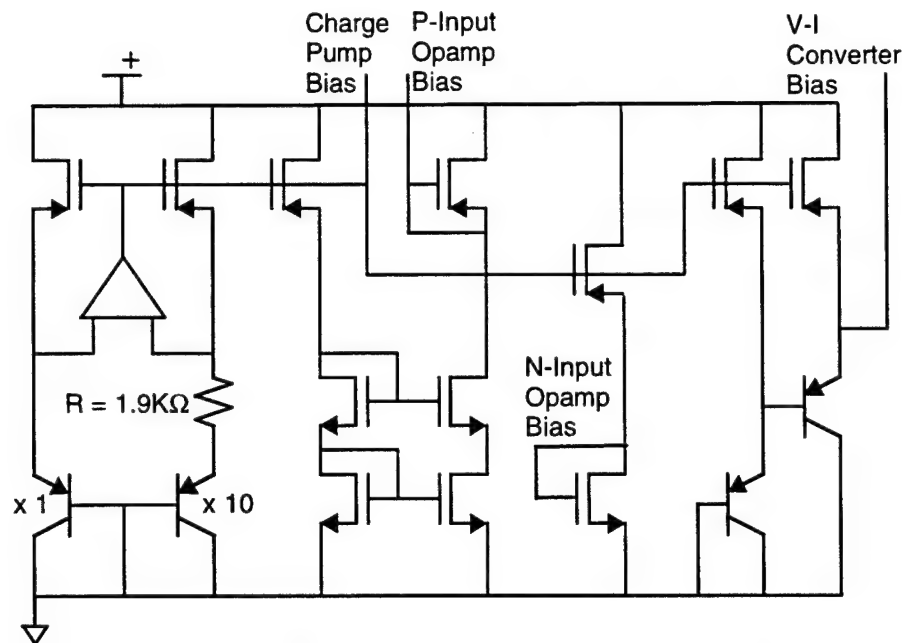


Figure 5.39 Left half of the bias generator schematic.

the required current bias for the three components mentioned above. Notice that the NMOS input opamp bias is obtained through appropriate current mirroring of the charge pump bias voltage. The bias voltage for the voltage-to-current converter block is set by the series connection of two forward-biased base-emitter junctions. This produces a voltage equal to $2V_{BE} = 1.4V$. It is also a very stable voltage as the exponential voltage to current relationship of the BJTs change V_{BE} little even for large changes in bias current. Since the bias current is derived from the kT/q generator, it is also very stable.

The right side of the circuit, shown in Figure 5.40, is essentially more of the same. The transistor ratio is set at 16, with a resistor value of $1.4 K\Omega$ to generate a current of $51.1 \mu A$. The PMOS devices are appropriately sized to scale the current to the desired $80 \mu A$ at the phase-frequency detector. The remainder of this circuit generates reference voltages in a manner similar to the V-I converter bias voltage discussed above. The kT/q bias current is mirrored, and generates a voltage across the resistor R_x . This voltage is applied to the base of a BJT, also biased at the kT/q current level, producing a voltage at the transistor emitter which is V_{BE} above the voltage across the resistor. The emitter

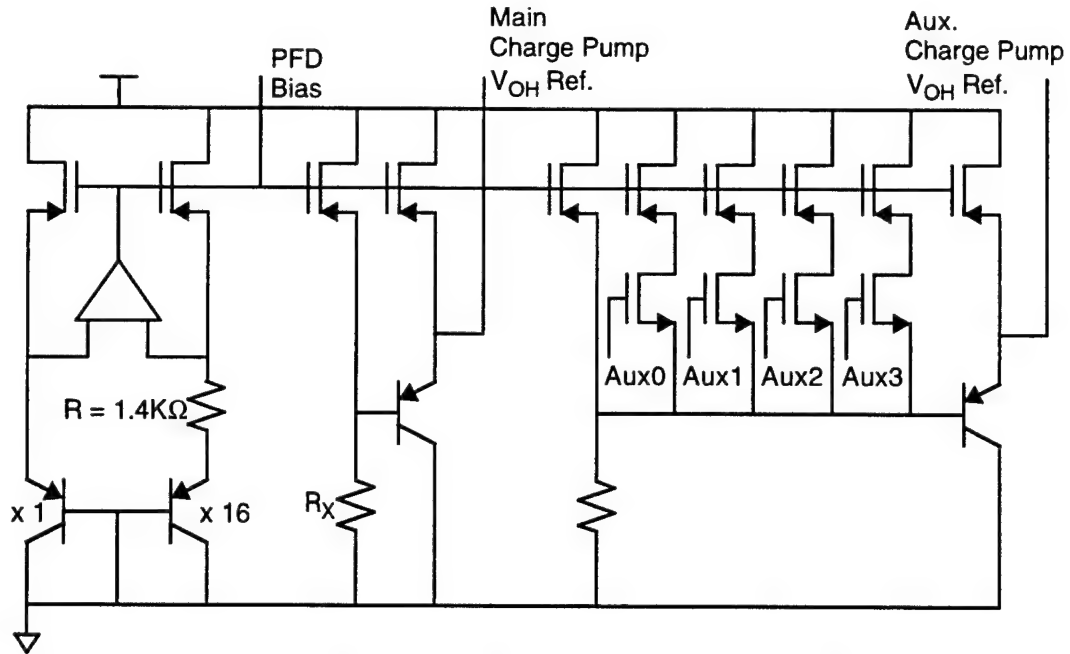


Figure 5.40 Right half of the bias generator schematic.

voltage is used as the main V_{OH} reference in the charge pump. In the generation of the auxiliary charge pump V_{OH} reference, four NMOS devices, gated by select signals, control the current that flows through the resistor. This provides a means of varying the V_{OH} level in four steps. Thus, four distinct values of A_f are achievable.

The plots of Figure 5.41 and Figure 5.42 demonstrate the bias generator's stability with respect to power supply voltage. Each bias current and reference voltage is displayed over a power supply voltage range of 1.8 - 2.2 volts. One should note that power supply sensitivity is worse for bias voltages far from that of the original kT/q generator. This is due to the fact that the simple current mirrors do not have the active opamp to improve

their r_o and, in turn, their current matching ability. This observation validates the decision to implement two separate kT/q generators.




Figure 5.41 Bias generator currents over various power supply voltages.




Figure 5.42 Bias generator voltages over various power supply voltages.

5.4 Simulation and Test Results

This section presents simulation and measurement results for the CMOS phase-locked loop clock generator. All simulation results given in this section were obtained using Berkeley SPICE bsim3 (version 3.1) models characterized from the wafer lot on which the PLL was run. These models exhibit an NMOS threshold voltage of 0.74, a 10% increase above the average value, and 5% increase over the maximum value used for simulations during the design phase.

Figure 5.43 shows the die photo of the CMOS PLL clock generator test chip. The chip consists of three versions of the phase-locked loop and several test structures. The clock generator layout measures $880\mu\text{m} \times 950\mu\text{m}$, with $608\mu\text{m} \times 608\mu\text{m}$ of that comprising the low-pass filter capacitor. This capacitor was created using the linear capacitor option of the HP14B process. It consists of poly-to-active capacitance, with the active in a well doped such that the structure exhibits a linear capacitance-to-voltage characteristic. The capacitor is composed of many unit-cell capacitors. These unit-cell capacitors, approximately 1pF in value, are also used as decoupling capacitors near sensitive circuitry, such as the voltage-controlled oscillator. The I/O pad drivers are sourced with a separate ring power and ground to isolate them from the PLL circuitry.




Figure 5.43 PLL Clock generator die photo.

Each PLL structure also has its own power and ground distribution. The pad drivers are standard, tapered-buffer structures with an open-drain output stage designed to drive a 50 Ω -to-ground output termination. As will be discussed shortly, the pad drivers caused a complication in the testing process.

5.4.1 CSA Phase Jitter Simulation

The inherent phase jitter of the CSA ring oscillator, predicted by equation (35), can be simulated using the method developed in Section 3.3. The results of the simulation demonstrate the oscillator's sensitivity to the intrinsic transistor noise of its components.

The jitter simulation begins with a noise analysis which provides the spectral noise density of the CSA oscillator with interstage interaction. In this simulation, the first stage is biased with a 0.8 V input, which is the midpoint of the CSA's voltage swing. The




Figure 5.44 Output noise spectral density for the CSA delay stage.




Figure 5.45 RMS Noise voltage for the CSA delay stage.

output of the second stage exhibits the voltage noise spectral density shown in Figure 5.44.

Following the example presented in Section 3.3, the noise spectral density must be integrated over the bandwidth of the CSA delay stage. The result of this analysis specifies the bounds over which the spectral noise density must be integrated. Figure 5.45 is a plot of the RMS noise voltage calculated by performing this integration and a subsequent square root operation. The marker represents the RMS noise voltage evaluated at the bandwidth of the circuit with $I_0 = 350\mu\text{A}$. At this frequency, the RMS noise voltage is $172\mu\text{V}$.

The RMS noise voltage represents the magnitude of the white noise generators which are applied to the ring oscillator in the transient simulation. With the noise




Figure 5.46 Simulated frequency spectrum of the CSA delay stage

generators between each stage of the ring oscillator, the transient simulation produces a waveform with a time varying frequency. The spread of the frequency distribution reflects the peak cycle-to-cycle phase jitter possible as a result of transistor noise. The resulting frequency spectrum is illustrated in Figure 5.46. The 0.6 MHz of frequency range, centered about 646.2 MHz, results in a simulated phase jitter of 1.4 ps. This is approximately twice the 0.73 ps predicted by the analytical method.

Figure 5.4 showed the PLL bandwidth to be approximately 400 KHz. This bandwidth results in an α (PLL phase jitter accumulation factor) of approximately 16. With $\alpha = 16$, the contribution of phase jitter due to inherent transistor noise in a CSA ring oscillator is 22.4 ps.

5.4.2 Measurement Results

The PLL clock generator measurements were taken by probing bare die with a high-frequency probe card. A Tektronix 11801 high-frequency digital sampling oscilloscope was used to observe the circuit's functionality.

The first measurement taken was an open loop frequency sweep of the PLL voltage-to-current converter and current-controlled oscillator. This test resulted in the voltage-to-frequency characteristic depicted in Figure 5.47.

The simulated trace represents data taken using the bsim3 models extracted from the wafer lot characterization data. The simulation predicts that the VCO will operate at a lower frequency than predicted by the initial simulations. This is due to the fact that the threshold voltages for the wafer run were 10% greater than the average values used in the initial simulations, and 5% greater than the maximum values used to verify PLL operation across process variation. The measured results proved to be slower yet. This is, most likely, due to increased junction and interconnect capacitance over that predicted by the




Figure 5.47 Voltage-to-frequency characteristic of open loop PLL.

layout extraction and circuit simulation. Increased capacitance would explain the near constant difference between the two curves.

Note also that the measured results do not reach as high a control voltage as the simulation. This is due to the aforementioned complication involving the pad drivers. With a pad ring V_{DD} of 3.3 V, the CSA circuits were unable to achieve a valid input-high level for the pad driver. This resulted in an apparent lack of activity on the oscillator output pin. Simulations showed that lowering the ring V_{DD} voltage to 1.8 V would bring the pad driver switching point low enough that the CSA circuits could drive it. Unfortunately, this fix had two side effects. First, the resulting output waveforms exhibit a low duty cycle, as illustrated in Figure 5.48. This is due to the fact that the CSA midpoint voltage does not coincide with the pad driver switching point. Second, the reduced ring V_{DD} prevents the pad driver from propagating the high end of the frequency range to the output pad. This prevents measuring the oscillator output frequency at the high end of the control voltage range, as shown in Figure 5.47.

The Tektronix oscilloscope simplifies the phase jitter measurements by incorporating the post-processing routines into the scope itself. The oscilloscope is




Figure 5.48 PLL output waveform.

capable of producing a histogram output illustrating the phase jitter of the measured signal. Inclusion of such oscilloscope traces is pending a means to obtain a hardcopy from the oscilloscope. However, the measured phase jitter results for the CMOS PLL clock generator are as follows. The peak-to-peak phase jitter, representing the absolute spread, was measured at 66.8 ps. The oscilloscope also determines the RMS cycle-to-cycle phase jitter, which was measured at 10.05 ps. Thus, the CMOS PLL clock generator design proved to be very stable in regards to phase jitter.

As a comparison, Table 3.1 is repeated here with an additional column representing the simulated and measured results of the clock generator reported in this work. As the data shows, the low-voltage CMOS clock generator reported in this work exhibits less than half the peak-to-peak phase jitter of the design reported in [3]. The results compiled here also correspond well to the measured peak-to-peak phase jitter results.

Table 5.3 Measured and simulated phase jitter results

	P-P Phase Jitter (ps) [3]	P-P Phase Jitter (ps) [this work]
Jitter contributor without supply noise		
White Noise in VCO	30	22.4
Dead zone of PFD	<10	0
Leakage on LF and Charge injection	15	0
Total Jitter without supply noise	55	22.4
Jitter due to a 0.2 V supply jump in 30 ps		
VCO induced jitter	80	24
Jitter induced by the change of the LF voltage	10	12.8
Total Jitter due to a 0.2 V supply jump	90	
Jitter due to a 10 mV substrate jump in 30 ps		
VCO induced jitter	<5	2.4
Total Jitter due to a 10 mV substrate jump	5	
Total Jitter (sum of the above contributors)	150	61.6

This chapter has detailed the design of a low-voltage, low phase jitter, phase-locked loop clock generator implemented in HP's CMOS14B process. The design

introduced and adapted a circuit called a current-steering amplifier which proved to be very versatile and effective in a low-voltage application. Nearly every one of the PLL components included the CSA. The charge pump implementation, in particular, achieved superior performance. The implementation completely eliminated the charge-sharing parasitic effect, and reduced the charge injection effect by 95% over typical rail-to-rail designs. Furthermore, the charge pump exhibited a very low steady-state error in current matching of 0.0478%.

The measured results presented in this chapter demonstrate the design's high-frequency operation at low power supply voltages. The design achieves nearly 700 MHz operation at 1.8 V in a 0.5 μm CMOS process. The phase-locked loop output clock signal was characterized by a peak-to-peak phase jitter of 66.8 ps, and an RMS cycle-to-cycle phase jitter of 10.05 ps.

The power supply voltage specification is very aggressive, given the threshold voltages for the 0.5 μm process. If full use of the process (5V) had been allowed, the clock generator design may have exhibited even better performance. The increased power supply voltage would extend the voltage headroom, driving the current source transistors further into saturation which provides increased tolerance to power supply variation.

CHAPTER VI

DELAY-LOCKED LOOP CLOCK GENERATION

This chapter explores the use of delay locked loops as clock generators.

Traditionally, delay-locked loops are not used in frequency synthesis applications because they require a frequency multiplication step that is not readily realized. The block diagram in Figure 6.1 demonstrates how a DLL adapts for use in clock generation. Very similar to a PLL, the DLL's negative feedback loop drives the input phase difference towards zero. However, the DLL differs significantly from a PLL in how it achieves this operation. While a PLL generates the signal which it phase-locks to the reference input, the DLL simply uses a delayed version of the input signal.

As in a PLL clock generator, there is a need for clock multiplication. In a digital circuit, frequency multiplication is most readily achieved through the EXOR logic function. Whether implemented as the common EXOR logic gate or a Gilbert cell multiplier, the EXOR function produces a logic high whenever its inputs are logically different (assuming a 2-input gate). If the two inputs are 90 degrees out of phase, or quadrature, the EXOR gate outputs the signal depicted in Figure 6.2. As the diagram illustrates, the output frequency is twice that of the input frequency. An added benefit to this technique is that the quadrature inputs produce an output signal which is also very

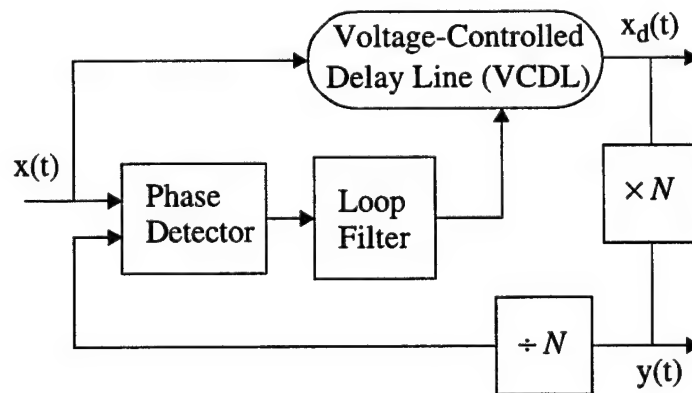


Figure 6.1 Delay-locked loop clock generator block diagram.




Figure 6.2 EXOR logic waveforms illustrating frequency multiplication.

nearly 50% in duty cycle.

Quadrature signals are readily available when the length of the delay line is a multiple of four. Steady state operation results in a VCDL delay of $M \cdot T_o$, an integer multiple of the input signal period. Most commonly the delay is a single input clock period [8]. In the locked state, each stage in the delay chain contributes $\frac{1}{N} \cdot T_o$ delay to the chain. Given that N is a multiple of 4, it is apparent that the output of each $\frac{N}{4}$ th delay stage is 90 degrees out of phase. The timing diagrams of Figure 6.3 illustrate this case for an 8 stage delay chain. The first 8 signals represent those at the output of each stage in the delay chain. The figure also illustrates the result of mixing the various outputs. Each successive level of multiplication provides an additional factor of 2 in frequency. Furthermore, multiplying quadrature signals in one level results in signals which are quadrature as well, as illustrated in Figure 6.3. Quadrature signals are available for $\log_2 N - 1$ levels, resulting in a maximum multiplication factor of $\frac{N}{2}$.

A concern with multiplying these signals is the potential for phase jitter as a result of power supply noise, especially considering the successive levels of signal multiplication. This concern can be mitigated through the use of the techniques discussed previously. Noise tolerant circuits, such as a Gilbert cell multiplier (differential logic), or a current-steering amplifier EXOR gate, reject noise on the power supply. Figure 6.4

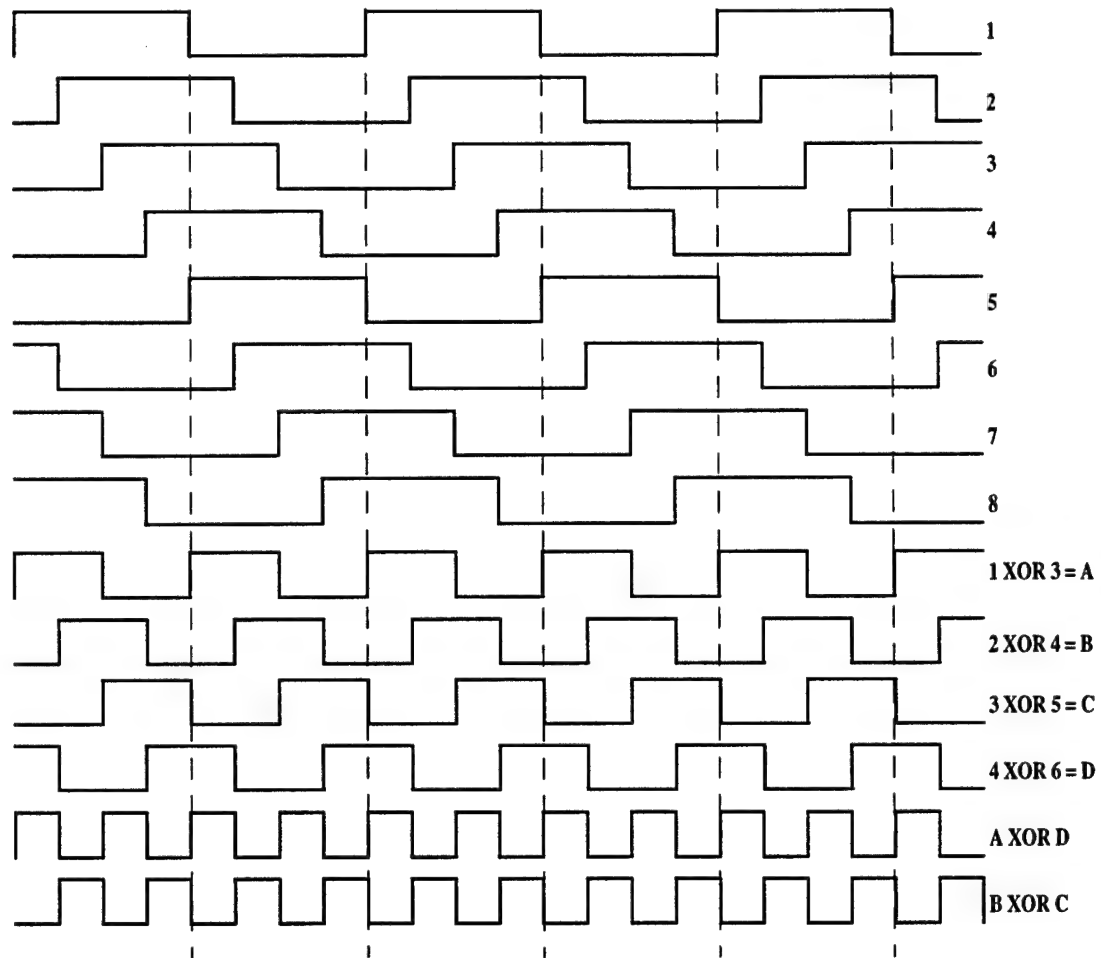


Figure 6.3 Logic waveforms in a DLL clock generator.

displays the circuit schematic for a Gilbert cell multiplier. Common complementary EXOR logic gates can be used if the power supply is clean or a clean reference voltage such as that described earlier from [19] can be used.




Figure 6.4 CMOS Gilbert cell multiplier.

The question of phase jitter also arises with regard to the voltage-controlled delay line. The analysis presented in [16], and repeated in this work, uses the first crossing approximation to estimate the phase jitter from inherent transistor noise in a specific delay stage. The phase jitter for both the source-coupled differential pair and the current steering amplifier have been presented in this work. The delay stage phase jitter is the same whether the stage is implemented in a voltage-controlled delay line or a voltage-controlled oscillator, but the phase jitter of these modules is very different.

In a ring oscillator the phase-jitter compounds cycle after cycle. Each perturbation that changes the delay of a stage away from the nominal causes the next stage to switch earlier, or later, as well. Each successive perturbation compounds upon the first. When the ring oscillator is incorporated into a phase-locked loop, the corrective action of the loop eliminates some of the compounded phase jitter. How much of this phase jitter is eliminated depends upon the bandwidth of the PLL. A fast reacting loop waits fewer cycles before correcting the phase jitter. A loop with a small bandwidth, however, allows more cycles to pass before corrective action is taken. Thus, the phase jitter of a VCO is multiplied by a factor α which is inversely proportional to the loop bandwidth [18].

The situation changes when the delay stage is incorporated into a voltage-controlled delay line in a delay-locked loop. When phase jitter occurs in a delay line, it propagates to the end of the delay chain and ends. There is no compounding of the error on successive cycles. While the phase jitter of a VCO in a PLL is multiplied by α , in a DLL $\alpha = 1$. Typical values of α for a PLL range from 10 to 100 [16].

It appears that if one is able to provide the frequency multiplication required for clock generation, a DLL should exhibit less overall phase jitter than a PLL. However, the comparison warrants a closer look. The high frequency VCO's required for next generation microprocessor clock generators rarely have more than three stages. A delay line to be used in a clock generator must have a number of stages equal to twice the highest required multiplication factor, unless other means of providing a 50% duty cycle

output are employed. Typical multiplication factors range from 4 to 8, so a delay line can easily be as many as 16 stages long.

The phase jitter of both a delay line and a VCO depend on the number of stages. Equations (11) and (34) represent the VCO phase jitter for the source-coupled differential pair and current-steering amplifier delay stages, respectively. Note that if these delay stages are used in a VCDL, the phase jitter will be 1/2 that predicted by these expressions. This is due to the fact that the derivation of timing jitter for an oscillator is proportional to the oscillation period which is given by $T_o = 2N \times t_d$ for a ring oscillator, and $T_o = N \times t_d$ for a delay line.

The phase jitter of these two specific delay stages also depends on a number of other factors including bias current, voltage swing, and voltage gain. However, many of these terms are interrelated. The phase jitter is inversely proportional to the bias current. Since the delay stage in a VCDL does not need to switch as quickly, and there are simply more delay stages overall, a lower bias current would offset the increase in power due to the increased number of stages. This decrease in bias current results in increased phase jitter. Thus, it is unclear whether or not the use of a DLL can improve phase jitter performance.

The picture is clarified significantly by expressing the timing variance as a function of only two variables: N and bias current. Starting with the source-coupled pair, this analysis proceeds as follows.

$$\overline{\Delta\tau_N^2} = \frac{kT}{I_{SS}} \frac{a_v \xi^2}{V_{GS} - V_T} \times T_o \quad (37)$$

Repeating equation (11) here for convenience, the first step is to introduce the dependence on N by replacing T_o with $T_o = 2N \times t_d$, where $t_d = V_{PP} \frac{C_L}{I_{SS}}$. This yields the expression

$$\overline{\Delta\tau_N^2} = \frac{kT}{I_{SS}^2} \frac{a_v \xi^2}{V_{GS} - V_T} \times 2NV_{PP}C_L \quad (38)$$

The components of this equation that are dependent upon supply current are a_v , $V_{GS}-V_T$, V_{pp} , and ξ .

$$a_v \equiv \frac{v_o}{v_{id}} = g_m \frac{r_o}{2} = \frac{V_A}{V_{GS}-V_T} \quad (39)$$

$$\frac{I_{ss}}{2} = K_n(V_{GS}-V_T)^2 \Rightarrow V_{GS}-V_T = \sqrt{\frac{I_{ss}}{2K_n}} \quad (40)$$

The voltage swing, represented by V_{pp} , can be approximated by the input voltage difference necessary to drive the differential pair into the unbalanced state.

$$\frac{V_{pp}}{2} \approx \sqrt{2}(V_{GS}-V_T) = \sqrt{\frac{I_{ss}}{K_n}}$$

Approximating the noise contribution factor, ξ , by the time-invariant version results in the following dependence upon bias current.

$$\xi = \sqrt{1 + \frac{2}{3}a_v} = \sqrt{1 + \frac{4}{3}V_A \frac{\sqrt{K_n}}{\sqrt{I_{ss}}}} \quad (41)$$

Substituting equations (39) through (41) into equation (38) yields the following expression for the time variance of the source-coupled pair.

$$\overline{v_N^2} = \frac{8kTC_L V_A N K_n}{I_{ss}^2 \sqrt{I_{ss}}} \cdot \left(1 + \frac{4}{3} V_A \frac{\sqrt{K_n}}{\sqrt{I_{ss}}} \right) \quad (42)$$

Given typical values, the quantity $\frac{4}{3} V_A \frac{\sqrt{K_n}}{\sqrt{I_{ss}}} > 1$ and equation (42) can be approximated by

$$\overline{v_N^2} \approx \frac{8kTC_L V_A N K_n}{I_{ss}^2 \sqrt{I_{ss}}} \cdot \left(\frac{4}{3} V_A \frac{\sqrt{K_n}}{\sqrt{I_{ss}}} \right) = \frac{32}{3} kTC_L V_A^2 K_n^{\frac{3}{2}} \frac{N}{I_{ss}^3} = D \frac{N}{I_{ss}^3} \quad (43)$$

The resulting expression shows that the timing variance actually has an inverse relationship to the cube of the bias current. This exposes an even higher sensitivity to

changes in bias current, and reinforces the trade-off between power dissipation and phase jitter performance.

Sections 3.2 and 5.3.3 showed that the time variance of a ring oscillator is implementation specific. Thus, it would be expected that different delay stages would exhibit different timing variance sensitivities to bias current. The timing variance of the current-steering amplifier delay stage's timing variance was derived in Section 5.3.3. A similar analysis reveals the relation of timing variance to bias current in equation (35), which is repeated here for convenience.

$$\overline{\Delta\tau_N^2} = \frac{4kT}{I_0} (1 + a_v^2) \xi^2 \frac{T_0}{\Delta V} \quad (44)$$

Once again, the first step is to replace T_0 with $T_0 = 2N \times t_d$, where $t_d = \frac{1}{2} \Delta V \frac{C_L}{I_0}$. This yields the expression

$$\overline{\Delta\tau_N^2} = \frac{4kTC_L}{I_0^2} (1 + a_v^2) \xi^2 N. \quad (45)$$

In this case, only a_v and ξ are functions of bias current. It is important to remember that the timing variance derivation for the current-steering amplifier assumed that the input voltage was equal to the output voltage. This bias condition represents the switching point of the gate where it is most susceptible to voltage noise and phase jitter.

$$a_v = g_{m1} R_{out} = g_{m1} \left(r_{o1} \parallel \frac{1}{g_{m2}} \parallel r_{o3} \right) \cong \frac{g_{m1}}{g_{m2}} \quad (46)$$

Since $V_{GS} - V_T$ is identical for both the driver and load transistors of the CSA stage (assuming that $V_{in} = V_{out}$), $\frac{g_{m1}}{g_{m2}}$ reduces to $\frac{(W/L)_1}{(W/L)_2}$.

The noise contribution factor, ξ , can be similarly reduced. Starting with the expression derived in Section 5.3.3, the relation changes as follows.

$$\xi^2 = \frac{1}{3} R_{out} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4} \right)^2 g_{m4} \right) = \frac{1}{3} \frac{g_{m1}}{g_{m2}} \left(g_{m1} + g_{m2} + g_{m3} + \left(\frac{W_3}{W_4} \right)^2 g_{m4} \right)$$

$$\xi^2 = \frac{1}{3} \frac{(W/L)_1}{(W/L)_2} \left(2 \sqrt{(K_{n1} + K_{n2})I_0} + 2 \sqrt{K_{n3}I_0} + \left(\frac{W_3}{W_4} \right)^2 2 \sqrt{K_{n4} \frac{W_4}{W_3} I_0} \right)$$

$$\xi^2 = A \sqrt{I_0}, \text{ where } A \text{ represents those values constant with respect to } I_0. \quad (47)$$

As a final step, equations (46) and (47) are substituted into equation (45). This step results in the expression for the CSA ring oscillator timing variance as given by equation (48).

$$\overline{\Delta \tau_N^2} = 4kTC_L \left(1 + \left[\frac{(W/L)_1}{(W/L)_2} \right]^2 \right) A \frac{N}{I_0 \sqrt{I_0}} = D \frac{N}{I_0 \sqrt{I_0}} \quad (48)$$

In equation (48), D represents the terms constant with respect to N and I_0 .

Equation (48) shows that the timing variance has a 3/2 power dependence on bias current. This is significantly less than the source-coupled differential pair. The result implies that the lower current required by a DLL delay stage would have a less significant impact on the overall phase jitter of the delay line were it implemented with CSA stages, rather than source-coupled pair stages.

To compare the phase jitter of a phase-locked loop with that of a delay-locked loop, the relationships represented by equations (43) and (48) can be evaluated over a range of operating parameters. Absolute values are avoided by normalizing the results to the phase jitter of a PLL with $N = 3$ and a bias current of 1 mA. The accumulation factor, α , is approximated by assuming an input frequency of 100 MHz. This implies that a bandwidth of 1 MHz is sufficient to render insignificant the amount of clock feedthrough present. A typical PLL with $BW = 10$ KHz exhibits a value of $\alpha = 100$, while a $BW = 100$ KHz exhibits a value of $\alpha = 30$ [16]. Since α is inversely proportional to the square root




Figure 6.5 Normalized phase jitter for source-coupled pair implementations.

of the PLL bandwidth, it should be readily apparent that a bandwidth of 1 MHz will result in an accumulation factor of 10.

This assumption allows the relations of equations (43) and (48) to be evaluated for PLL's and DLL's where $\alpha = 10$ for the PLL and $\alpha = 1$ for the DLL. Figure 6.5 depicts the normalized phase jitter of a PLL with $N = 3$, and three DLL's with N values of 8, 16, and 32. The vertical axis represents the evaluation of equation (43), divided by the phase jitter of a PLL with a 3-stage oscillator and operating at a bias current of 1 mA. Thus, the curve for the PLL intersects the $Y = 1$ line at a bias current value of 1 mA.

As the figure illustrates, the cubic dependence of phase jitter on bias current causes the phase jitter to be significantly higher at low bias currents. This is particularly detrimental to the DLL which must operate at lower bias currents due to the increased number of stages, and the need for a longer delay through the chain. Figure 6.6 expands

the view by plotting only the normalized phase jitter of the delay-locked loops with the line at $Y = 1$ depicting the nominal PLL phase jitter performance.

The DLL phase jitter is significantly higher than PLL phase jitter in the region of interest. Since the DLL operates at the same frequency as the PLL input reference signal, and the delay chain typically contains significantly more stages, the required bias current lies on the low end of the scale. Figure 6.6 shows that this results in a phase jitter that is, at best, comparable to that of a PLL. This occurs despite the fact that the PLL accumulates excess phase jitter as a function of α .

When one considers the same relationship for implementations using the current-steering amplifier delay stage, the result is very different. Equation (48) reveals that the phase jitter of a CSA oscillator is dependent upon the bias current, but at a significantly lower sensitivity. The inverse $3/2$ power contribution of bias current to the phase jitter provides a much more gradual increase of phase jitter at lower bias current levels.




Figure 6.6 Normalized phase jitter of the DLL in comparison to the PLL

The normalized phase jitter predicted by equation (48) is plotted in Figure 6.7 versus bias current for both a PLL and DLL; the PLL has $N = 3$. As in the case of the source-coupled pair delay stage, the phase jitter is normalized to that of a PLL with $N = 3$ and $I = 1\text{mA}$.

Two differences to the source-coupled pair implementation stand out immediately. First, the peak of the normalized phase jitter at low bias current levels is much less than that of the previous case. Due to the weaker dependence of the phase jitter on bias current, an implementation using the CSA delay stage can operate at a lower frequency without compromising the phase jitter performance. Second, the phase jitter of the PLL is significantly higher than that of the various DLLs, even at high current levels.




Figure 6.7 Normalized phase jitter for current-steering amplifier implementations.

This observation is depicted more clearly in Figure 6.8 which expands the view by plotting only the normalized phase jitter of the delay-locked loops with the line at $Y = 1$ depicting the nominal PLL phase jitter performance.

As the plots illustrate, only at the lowest bias current levels is the phase jitter performance of the DLL worse than that of the PLL. While this makes no statement about the relative phase jitter magnitude of the CSA versus the source-coupled pair, it does show that for situations in which the source-coupled pair is not an option, such as low-voltage applications, the DLL is a better clock generator than the PLL.




Figure 6.8 Normalized phase jitter of the DLL in comparison to the PLL.

CHAPTER VII

CONCLUSIONS

Market pressure to continually increase microprocessor frequency has pushed designs into the range where clocking issues, such as clock skew and phase jitter, have become significant problems. These obstacles to next generation microprocessor design require new simulation methods and circuit topologies.

One area requiring such attention is microprocessor clock generation. A reduced clock error budget prompts the need for designs exhibiting solid frequency stability. Industry trends complicate this task by steadily reducing the power supply voltage. This work has explored phase jitter, developing both simulation methods and circuit topologies to minimize its effects in low-voltage, high frequency microprocessor designs. This work has led to several significant contributions, which are highlighted in the following sections.

7.1 Contributions

7.1.1 CGaAs PLL Clock Generator

The clock generator reported in this work represents the first such circuit designed and tested in Motorola's Complementary GaAs process. The 120 ps of phase jitter (measured as absolute phase jitter), 1.5 V power supply voltage, and 800 MHz maximum VCO frequency demonstrated the design's superior performance. Furthermore, the PLL remained operational at power supply voltages as low as 0.8 V, illustrating its viability for use in low voltage applications.

7.1.2 Phase Jitter

This work compiles the current knowledge of phase jitter, providing background information, analytical methods, examples, and design guidelines for low-phase jitter

clock generators. A phase jitter simulation methodology was developed, which provides a more accurate means of determining the oscillator phase jitter due to inherent transistor noise than the analytical method. It includes the phase jitter model in transient PLL simulations so that one can evaluate the effects of oscillator phase jitter on the PLL's tracking behavior. Furthermore, the methodology provides the framework for including additional noise effects into the oscillator model.

7.1.3 CMOS PLL Clock Generator

Using Hewlett Packard's 0.5 μm digital CMOS process, a CMOS PLL clock generator was designed which achieved nearly 800 MHz operation at power supply voltages as low as 1.8 V. This achievement was due, in large part, to the adaptation of the current-steering amplifier. This circuit served as the foundation for many of the PLL components. Its versatility and low voltage requirements suit it well to such designs. Of particular interest, the current-steering amplifier charge pump eliminates charge sharing, reduces charge injection by 95 % over some designs, and exhibits only 0.048 % error in its steady-state operation.

Simulation and measurement results also demonstrated the design's excellent phase jitter performance. Laboratory measurements show a peak-to-peak absolute phase jitter of 60 ps, and an RMS cycle-to-cycle phase jitter of 10 ps. This performance surpasses that of the PLL reported in [61], which achieved a cycle-to-cycle phase jitter of 12 ps, but utilized the voltage reference method described in Section 3.2 to provide a clean 2.0 V power supply from a noisy 3.3 V source.

7.1.4 Delay-Locked Loop Clock Generation

Traditionally, clock generators have been designed as phase-locked loops. This is largely due to the relative ease of frequency multiplication in phase-locked loops as

compared with delay-locked loops. However, this work has shown that the phase jitter performance of a delay-locked loop can surpass that of a similarly designed phase-locked loop. This was demonstrated through the derivation of the phase jitter dependence on bias current for various delay stage implementations.

7.2 Future Work

This work both contributed to the area of clock generator design and exposed areas which warrant additional research.

The current-steering amplifier has proven to be a useful and versatile circuit in low-voltage applications. The performance of the circuit could be improved by replacing the single PMOS current source transistor with a better controlled current source. This could be realized in the following ways. First, a current source network, such as a cascode or Wilson current source, could be used to improve the current matching ability and power supply noise rejection of the CSA. Also, since NMOS devices make better current sources than PMOS devices, the dual of the CSA circuit used in this work and illustrated in Figure 7.1 should be evaluated.

While the CSA demonstrated good power supply rejection properties, its ability to operate at low power supply voltages indicates that the combination of CSA circuits and a locally regulated power supply, such as that described in Section 3.2, would exhibit even




Figure 7.1 CSA dual circuit diagram

better performance in a noisy digital environment. The viability of such a combination in state of the art CMOS processes should be explored.

Chapter VI shows that a delay-locked loop implementation of a clock generator could exhibit lower phase jitter than one that is phase-locked. There is no disputing the fact that a delay-line is more stable than an oscillator. It remains to be seen if a full DLL clock generator implementation would result in better phase jitter performance with acceptable levels of power dissipation.

Finally, the phase jitter simulation methodology developed in this work applies only to the full circuit PLL model. This model requires long simulation times to adequately characterize the tracking behavior of the PLL. A very useful extension of this work would incorporate the phase jitter noise model into the behavioral model of the oscillator. This would allow efficient system-level simulations with a complete noise model.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] Alvarez et al., A Wide-Bandwidth Low-Voltage PLL for PowerPC Microprocessors *IEEE Journal of Solid-State Circuits*, pp. 383-391, April 1995.
- [2] J.G. Maneatis, "Low-Jitter Process-Independent DLL and PLL Based on Self-Biased Techniques," *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 11, pp. 1723-1732, November 1996.
- [3] V. von Kaenel, D. Aebischer, et. al., "A 320 MHz, 1.5 mW @ 1.35 V CMOS PLL for Microprocessor Clock Generation," *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 11, pp. 1715-1722, November 1996.
- [4] S. Stetson and R. Brown, "A Complementary GaAs PLL Clock Multiplier with Wide-Bandwidth and Low-Voltage Operation," *IEEE GaAs IC Symposium Technical Digest*, pp. 317-320, 1996.
- [5] K. Lalgudi, and M. Papaefthymiou, "Retiming edge-triggered circuits under general delay models," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 16, No. 12, pp. 1393-1408, December 1997.
- [6] A. Wolfe, "Internal layout of Intel's Merced comes to light," *EE Times Online*, September 28, 1998.
- [7] R. E. Best, "Phase-Locked Loops: Theory, Design, and Applications." New York, NY: McGraw-Hill, 1993.
- [8] B. Razavi, "Design of Monolithic Phase-Locked Loops and Clock Recovery Circuits - A Tutorial," *Monolithic Phase-Locked Loops and Clock Recovery Circuits: Theory and Design*, Piscataway, NJ, IEEE Press, 1996.
- [9] F.M. Gardner, "Charge-Pump Phase-Locked Loops", *IEEE Transactions on Communications*, Vol. COM-28, pp. 1849-1858, November, 1980.
- [10] A. Pottbacker and U. Langmann, "An 8 GHz silicon bipolar clock recovery and data regenerator IC," *IEEE Journal of Solid-State Circuits*, Vol. 29, pp. 1572-1576, December 1994.
- [11] B. Razavi and J. Sung, "A 2.5-Gb/sec 15-mW BiCMOS clock recovery circuit," *Symposium of VLSI Circuits Digest of Technical Papers*, pp. 183-185, 1985.
- [12] E. Friedman, "Clock Distribution Networks in VLSI Circuits and Systems." New York, NY: IEEE Press, 1995.
- [13] D. Wann and M. Franklin, "Asynchronous and Clocked Control Structures for VLSI Based Interconnection Networks," *IEEE Transactions on Computers*, Vol. C-32, No. 3, pp. 284-293, March 1983.

- [14] V. F. Kroupa, "Noise properties of PLL systems," *IEEE Transactions on Communications*, Vol. COM-30, pp. 2244-2252, October 1982.
- [15] I.A. Young, et. al., "A PLL Clock Generator with 5 to 110 MHz of Lock Range for Microprocessors," *IEEE Journal of Solid-State Circuits*, Vol. 27, No. 11, pp. 1599-1607, November 1992.
- [16] T. Weigandt, B. Kim, P. Gray, "Analysis of Timing Jitter in CMOS Ring Oscillators," *ISCAS 1994 Proceedings*, 1994.
- [17] A. Abidi, R. Meyer, "Noise in Relaxation Oscillators," *IEEE Journal of Solid State Circuits*, vol. SC-18, December 1983.
- [18] B. Kim, T.C. Weigandt, P.R. Gray, "PLL/DLL System Noise Analysis for Low Jitter Clock Synthesizer Design," *ISCAS 94 Proceedings*, June 1994.
- [19] K. Ware, H. Lee, C. Sodini, "A 200-MHz CMOS Phase-Locked Loop with Dual Phase Detectors," *IEEE Journal of Solid-State Circuits*, vol. 24, pp. 1560-1568, December 1989.
- [20] I. Novof, "Fully Integrated CMOS Phase-Locked Loop with 15 to 240 MHz Locking Range and ± 50 ps Jitter," *ISSCC Digest of Technical Papers*, pp. 112-113, February 1995.
- [21] D. Mijuskovic et al., "Cell-based fully integrated CMOS frequency synthesizers," *IEEE Journal of Solid-State Circuits*, Vol. 29, pp. 271-279, March 1994.
- [22] *Saber Simulation Reference Manual*, Analogy Inc.
- [23] F.M. Gardner, *Phaselock Techniques*, 2nd ed. New York: Wiley, 1979.
- [24] F.M. Gardner, "Phase Accuracy of Charge Pump PLL's," *IEEE Transactions on Communications*, Vol. COM-30, pp. 2362-2363, October 1982.
- [25] M. Van Paemel, "Analysis of a Charge-Pump PLL: A New Model," *IEEE Transactions on Communications*, Vol. 42, pp. 2490-2498, July 1994.
- [26] D. Jeong, et. al., "Design of PLL-Based Clock Generation Circuits," *IEEE Journal of Solid-State Circuits*, Vol. SC-22, pp. 255-261, April 1987.
- [27] B. Razavi, "Analysis, Modeling, and Simulation of Phase Noise in Monolithic Voltage-Controlled Oscillators," *Proceedings of the Custom Integrated Circuits Conference*, May 1995.
- [28] S. Ohr, "Analog technologists decry plummeting circuit voltage," *EE Times*, pp. 75-78, July 13th, 1998.

- [29] D.J. Allstot, G. Liang, H.C. Yang, "Current-mode logic techniques for CMOS mixed-mode ASIC's," *Proceedings of IEEE Custom Integrated Circuits Conference*, pp. 25.2.1-25.2.4, 1991.
- [30] H.C. Yang, L.K. Lee, R.S. Co, "A Low Jitter 0.3-165 MHz CMOS PLL Frequency Synthesizer for 3 V/5 V Operation," *IEEE Journal of Solid-State Circuits*, Vol. 32, No. 4, pp. 582-586, April 1997.
- [31] W. Egan, "Phase Noise Modeling in Frequency Dividers," *Proceedings of the 45th Annual Symposium on Frequency Control*, pp. 629-635, 1991.
- [32] P. Gronowski, et. al., "A 433-MHz 64-b Quad-Issue RISC Microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 11, pp. 1687-1695, November 1996.
- [33] S.H. Unger and C-J. Tan, "Clocking Schemes for High-Speed Digital Systems," *IEEE Transactions on Computers*, Vol. C-35, No. 10, pp. 880-895, October 1986.
- [34] I. Lin, J.A. Ludwig, and K. Eng, "Analyzing Cycle Stealing on Synchronous Circuits with Level-Sensitive Latches," *Proceedings of ACM/IEEE Design Automation Conference*, pp. 393-398, June 1992.
- [35] R.-S. Tsay and I. Lin, "Robin Hood: A System Timing Verifier for Multi-Phase Level-Sensitive Clock Designs," *Proceedings of IEEE International Conference on ASICs*, pp. 516-519, September 1992.
- [36] H. B. Bakoglu, "Circuits, Interconnections, and Packaging for VLSI." New York, NY: Addison Wesley, 1990.
- [37] D. W. Dobberpuhl, et al., "A 200-MHz 64-b Dual Issue CMOS Microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. SC-27, No. 11, pp. 1555-1565, November 1992.
- [38] M. Horowitz, "Clocking Strategies in High Performance Processors," *Proceedings of the IEEE Symposium on VLSI Circuits*, pp. 50-53, June 1992.
- [39] J. McNeill, "Jitter in Ring Oscillators," *1994 IEEE International Symposium on Circuits and Systems*, pp. 201-204, vol.6, 1994.
- [40] M. Williams, "A Discussion of Methods for Measuring Low-Amplitude Jitter," *Proceedings of the International Test Conference*, pp. 646-652, 1995.
- [41] R. Co, J.H. Mulligan, "Optimization of Phase-Locked Loop Performance in Data Recovery Systems," *IEEE Journal of Solid-State Circuits*, Vol. 29, pp. 1022-1034, September 1994.
- [42] J. Montanaro, R. T. Witek, et. al., "A 160-MHz, 32-b, 0.5-W CMOS RISC

- Microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. 31, No. 11, pp. 1703-1714, November 1996.
- [43] E. De Man and M. Schobinger, "Power Dissipation in the Clock System of Highly Pipelined ULSI CMOS Circuits," *Proceedings of the International Workshop on Low Power Design*, pp. 133-138, April 1994.
 - [44] H. Kojima, S. Tanaka, and K. Sasaki, "Half-Swing Clocking Scheme for 75% Power Saving in Clocking Circuitry," *Proceedings of the IEEE Symposium on VLSI Circuits*, pp. 23-24, June 1994.
 - [45] Mihai Banu, CMOS Oscillators with Multi-Decade Tuning Range and Gigahertz Maximum Speed *IEEE Journal of Solid State Circuits*, pp.1386-1393, December, 1988.
 - [46] Deog-Kyoon Jeong et al., Design of PLL-Based Clock Generation Circuits *IEEE Journal of Solid-State Circuits*, pp. 255-261, April 1987.
 - [47] N. Weste and K. Eshraghian, "Principles of CMOS VLSI Design," 2nd ed. New York, NY: Addison-Wesley, 1993.
 - [48] L. Gwennap, "Digital Leads the Pack with 21164," *Microprocessor Report*, Vol. 8, No. 12, September 1994.
 - [49] B. Benschneider, et. al., "A 300-MHz 64-b Quad-Issue CMOS RISC Microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 11, pp. 1203-1211, November 1995.
 - [50] G. Di Cataldo, G. Palumbo, "New CMOS Schmitt Triggers," *ISCAS 1992 Proceedings*, 1992.
 - [51] Q. Zhu, W. Dai, "Planar Clock Routing for High Performance Chip and Package Co-Design," *IEEE Transactions on VLSI Systems*, Vol. 4, No. 2, pp. 210-226, June 1996.
 - [52] J. Neves, E. Friedman, "Circuit Synthesis of Clock Distribution Networks based on Non-Zero Clock Skew," *Proceedings of IEEE International Symposium on Circuits and Systems*, pp. 4.175-4.178, May/June 1994.
 - [53] L. Benini, G. De Micheli, "Transformation and synthesis of FSMs for low-power gated implementation," *Proceedings 1995 International Symposium on Low Power Design ACM*, pp. 21-26, April 1995.
 - [54] C. Nagendra, M.J. Irwin, "Design trade-offs in CMOS FIR filters," *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, Vol. 6, May 1996.
 - [55] L. Hall, et al., "Clock Distribution Using Cooperative Ring Oscillators,"

Proceedings of 1997 ARVLSI Conference, September 1997.

- [56] S. Pullela, N. Menezes, L.T. Pillage, "Reliable Non-Zero Clock Trees Using Wire Width Optimization," *Proceedings of ACM/IEEE Design Automation Conference*, pp. 165-190, June 1993.
- [57] H.B. Bakoglu, J.T. Walker, J.D. Meindl, "A Symmetric Clock Distribution Tree and Optimized High-Speed Interconnections for Reduced Clock Skew in ULSI and WSI Circuits," *Proceedings of IEEE International Conference on Computer Design*, pp. 118-122, October 1986.
- [58] B. Razavi, "A Study of Phase Noise in CMOS Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 31, March 1996.
- [59] P. Gray, R. Meyer, "Analysis and Design of Analog Integrated Circuits," John Wiley & Sons, NY, 1993.
- [60] A. Sedra, C. Smith, "Microelectronic Circuits," 3rd ed., Saunders College Publishing, Philadelphia, PA, 1991.
- [61] V. von Kaenel, et al., "A 600 MHz CMOS PLL Microprocessor Clock Generator with a 1.2 GHz VCO," *1998 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, pp. 396-397.
- [62] S. Sun, "An Analog PLL-Based Clock and Data Recovery Circuit with High Input Jitter Tolerance," *IEEE Journal of Solid-State Circuits*, Vol. 24, pp. 325-330, April 1989.
- [63] M. Wakayama, A. Abidi, "A 30-MHz Low-Jitter High-Linearity CMOS Voltage-Controlled Oscillator," *IEEE Journal of Solid-State Circuits*, Vol. SC-22, pp. 1074-1080, December 1987.
- [64] A. Thaik, H. Nguyen, "A Dual PLL Based Multi Frequency Clock Distribution Scheme," *1992 Symposium on VLSI Circuits Digest of Technical Papers*, pp. 84-85, 1992.
- [65] J. Vital, C. Temes, "Clock Generation System with Reduced Jitter Noise in the Baseband," *1991 IEEE International Symposium on Circuits and Systems*, pp. 2621-2624, 1991.
- [66] Z. Zhang, et. al., "A 360 MHz CMOS PLL with 1V Peak-to-Peak Power Supply Noise Tolerance," *1996 IEEE International Solid-State Circuits Conference*, pp. 134-135, 1996.
- [67] S. Kim, et. al., "A 960 Mbps/pin Interface for Skew-Tolerant Bus Using Low Jitter PLL," *1996 Symposium on VLSI Circuits Digest of Technical Papers*, pp. 118-119, 1996.

- [68] M. Johnson, "A Variable Delay Line PLL for CPU-Coprocessor Synchronization," *IEEE Journal of Solid-State Circuits*, Vol. 23, pp. 1218-1223, October 1988.
- [69] S. Sidiropoulos, M. Horowitz, "A Semi-Digital Dual Delay Locked Loop," *IEEE Journal of Solid State Circuits*, Vol.32, pp. 1683-1692, Nov. 1997.
- [70] A. Efendovich, et. al., "Multifrequency Zero-Jitter Delay-Locked Loop," *IEEE Journal of Solid-State Circuits*, Vol. 29, pp. 67-70, January 1994.
- [71] J. McNeill, R. Croughwell, "A 150 mW, 155 MHz Phase Locked Loop with Low Jitter VCO," *1994 IEEE International Symposium on Circuits and Systems*, pp. 49-52, vol.3, 1994.
- [72] D. Woeste, et. al., "Digital-Phase Aligner Macro for Clock Tree Compensation with 70ps Jitter," *1996 IEEE International Solid-State Circuits Conference*, pp. 136-137, 1996.
- [73] R. Khanna, et. al., "A 0.25mmx86 Microprocessor with a 100MHz Socket 7 Interface," *IEEE International Solid-State Circuits Conference*, pp. 242-243, 1998.
- [74] N. Rohrer, et. al., "A 480MHz RISC Microprocessor in a 0.12 μm L_{eff} CMOS Technology with Copper Interconnects," *IEEE International Solid-State Circuits Conference*, pp. 240-241, 1998.
- [75] J. Silberman, et. al., "A 1.0GHz Single-Issue 64b PowerPC Integer Processor," *IEEE International Solid-State Circuits Conference*, pp. 230-231, 1998.
- [76] H. Fair, D. Bailey, "Clocking Design and Analysis for a 600MHz Alpha Microprocessor," *IEEE International Solid-State Circuits Conference*, pp. 398-399, 1998.
- [77] G. Geannopoulos, X. Dai, "An Adaptive Digital Deskewing Circuit for Clock Distribution Networks," *IEEE International Solid-State Circuits Conference*, pp. 400-401, 1998.